

Lotteries, Queues, and Bottlenecks

Gil Hersch and Thomas Rowe

1. Introduction

Consider the choices those in charge of the US immigration system make when they determine to whom to grant an immigration visa, colloquially known as a green card. How should the US distribute immigration visas? A common approach to distributional questions is to argue that we should distribute goods based on some morally relevant criteria, such as desert or effort, and there is a vast literature discussing what these criteria are (Milne 1986; Miller 1989; Lamont 1997; Rawls (1999); Knight 2011; Fumagalli 2022). In the case of those seeking to immigrate to the US, there are a variety of morally relevant differences between applicants; whether they are refugees, have family members who are US citizens, or those who have “extraordinary ability.”¹ It is reasonable to claim that such differences between immigration visa applicants are morally relevant to the decision of whether to grant them a visa.²

Nevertheless, the overall demand for immigration visas to the US far exceeds the current supply.³ Assuming the US still has some visas left in its

¹ These are three categories that the US immigration system treats separately.

² Although it is reasonable to claim that such differences are morally relevant, we are not assuming the correctness of such a “threshold” approach that accords lexical priority to those who meet certain moral criteria before permitting discretionary immigration. For example, the Canadian points-based immigration system allows for a complete ranking of candidates according to various criteria that they deem morally relevant without according lexical priority. Nevertheless, there are many cases where every claimant has an equally strong claim to receive the good in question, whether because there are no morally relevant differences, or such differences were already addressed. Our project addresses these kinds of cases.

³ It is true that the US could simply allow more immigration than it currently does, but unless one is committed to open borders (not an unreasonable position), then it is commonly accepted that the US is entitled to its discretion with respect to how many visas it issues. For a discussion of open borders see, for example, Carens (1987) and Van der Vossen and Brennan (2018). Moreover, given that around thirteen million people applied for the US immigration visa lottery in 2020, even if the US increased the number of visas it supplies ten or a hundredfold, it would still not be enough to meet demand.

quotas to dispense to those with no special moral claims, what then? One way of distributing goods involves the decidedly economic solution of the price mechanism through markets. Much has been said in favor of this elegant solution.⁴ The price mechanism assigns each good a price that will clear the market. This equalizes supply and demand and deems any good non-scarce at that price level, because anyone who wants the good and is willing to pay that price can get it. There are, however, a variety of reasons to worry that such a solution often is unfair, and unfair in the specific case of immigration visas.⁵ Consider our attitudes to the myriad of cases in which countries grant the proverbial “gold visas”—visas that can be obtained in exchange for money. These often stoke a moral uproar regarding the lack of fairness in granting the rich yet another way by which their money talks. Hidalgo (2016) references several such reactions when arguing for the permissibility of the sale of immigration visas (Shachar 2009: 54–61; European Parliament 2013; Reding 2014; Shachar & Hirschl 2014a: 250; Shachar & Hirschl 2014b: 248). That one’s willingness to pay, no matter how great, should not grant them special consideration in cases such as immigration visas is a commonly held view.

If we refrain from relying on markets for immigration visas to prevent widespread outrage, what are we left with? This chapter focuses on the problem of distributing goods in cases where there are no morally relevant differences among those demanding the good (from the distributor’s perspective), and reliance on markets is considered morally inappropriate.

The two most commonly discussed alternatives, both generally and in the academic literature, are lotteries and queues. Lotteries are randomization procedures that assign potential recipients a chance of receiving a good. Queues are a ubiquitous distribution mechanism in which individuals enter an ordered temporal sequence to receive a good sequentially based on the order in which the individuals entered the queue. Queuing mechanisms are also sometimes referred to as “First Come, First Served” (FCFS) or “First-In-First-Out” (FIFO).

⁴ See for example Robbins (1935).

⁵ Authors that argue for the inappropriateness of the price mechanism for allocating certain goods include Sandel (2012) and Satz (2012). One easy to grasp worry is that in a society in which there is a background of unjust distribution (some are unjustly wealthier than others), using the price mechanism to distribute goods will usually result in continuing the injustice. We thank Brian Berkey for this comment. Moreover, the fact that price gouging is often portrayed in a negative light is a case in point. See, for example, Lamont and Favor (2009); Snyder (2009); Brake (2020).

But which distributive method is more appropriate to use in the case of immigration visas? To answer this question we introduce a third, novel, type of case in addition to cases of abundance and cases of scarcity—bottleneck cases—and argue that lotteries are more appropriate in cases of scarcity, whereas queues are more appropriate in bottleneck cases.⁶ An upshot of our discussion is that it entails that immigration visas to the US should be understood as a case of scarcity, and therefore it is the kind of case to which lotteries are more appropriate than queues.

This is indeed what the US government more or less does through its Diversity Immigrant Visa Program, which makes up to 50,000 immigrant visas available annually, drawn from random selection among all entries to individuals who are from countries with low rates of immigration to the United States.⁷ Given that demand for immigration visas to the US far outstrips the supply, we treat it as a case of scarcity and consequently we view the current practice of distributing this subset of visas through a lottery as appropriate.

More generally, we argue that both arguments for lotteries and for queues have merit, albeit for different distributive scenarios. Cases of scarcity arise when demand for the good cannot be met by the current supply.⁸ Attempting to fairly distribute waiting times when only a subset of those demanding the good will get it, entails that some of the group will, in effect, be waiting forever. However, a secondary good that can be distributed is the chances of obtaining the good. This secondary good becomes an option when it is not possible for everyone who needs the good to receive it. In cases

⁶ Issues of distributive justice do not arise in cases of abundance (Goodin 2001). David Hume also makes this clear when he rhetorically asks “For what purpose make a partition of goods, where everyone has already more than enough? [...] Justice, in that case, being totally useless, would be an idle ceremonial, and could never possibly have place in the catalogue of virtues” (Hume 1777).

Hume also argues that considerations of justice do not arise in extreme scarcity:

Suppose a society to fall into such want of all common necessities, that the utmost frugality and industry cannot preserve the greater number from perishing, and the whole from extreme misery; it will readily, I believe, be admitted, that the strict laws of justice are suspended, in such a pressing emergence, and give place to the stronger motives of necessity and self-preservation. (114)

We leave open whether Hume is right in this regard as it should not affect our argument in this article.

⁷ <https://www.uscis.gov/green-card/green-card-eligibility/green-card-through-the-diversity-immigrant-visa-program>.

⁸ We make no reference to any difference between needs and wants, since we view this distinction both highly contentious as well as immaterial for our purposes. We view a broader term—demand—as encompassing both without making a distinction between them.

of scarcity, we submit, lotteries are the most appropriate way to allocate the good. When dealing with bottlenecks, everyone will eventually get the good, so equalizing chances of obtaining the good is meaningless. Instead, the secondary good that can and ought to be distributed fairly is waiting time. Queues are the most appropriate way to distribute the good of waiting time. In what follows, we precisify the concepts of abundance, scarcity, and bottlenecks before demonstrating why the presence of each in a distributive scenario warrants a different response.

2. Abundance, Scarcity, and Bottlenecks

We begin by outlining the concepts of abundance, scarcity, and bottlenecks. We consider a case of abundance as one in which demand for the good is at least as high as supply at a given time. Goods like air, which is free, but also goods that cost money but are fully stocked, like (pre-pandemic) toilet paper, or readily available, like tap water, are goods we consider to be abundant.

To illustrate this point, consider *Camping on BLM Land*.

The Bureau of Land Management (BLM) is an agency within the United States Department of the Interior responsible for administering federal lands. Most of the public lands away from developed recreation facilities are open to dispersed camping, without any permits or fees required. There is always ample space for anyone interested in camping on BLM land, especially in the Western US. There are, given the demand, no costs, no bureaucratic hurdles, no limits on occupancy, and no limits on how and when anyone enters. Camping on BLM land is an abundant good.

We consider cases of scarcity to arise when the demand for the good at a given time cannot be fully satisfied.⁹ Positional goods are an ideal example of scarcity since such goods' value arises, by definition, from them not being obtainable to everyone who might want them. No matter how much time goes by, only some people will have the fastest car, the fanciest house, or the

⁹ To this extent, we are adopting the definition in economics of "relative scarcity," a condition where there is not enough of a resource to satisfy existing demand, rather than "absolute scarcity," which refers to resources that are limited regardless of demand (Raiklin and Uyar 1996).

best grades. Other examples of scarcity include cases in which there are not enough kidneys for all those who must suffer life on dialysis and early death if they go without, green cards when there are more would-be immigrants than a country will allow, as well as more mundane cases like tickets to the World Cup finals or overhead storage bins in an already cramped airplane.

To illustrate, consider *Rafting in the Grand Canyon*.¹⁰

The stretch of the Colorado River that runs through the Grand Canyon is one of the most popular white-water rafting destinations in the US. In order to protect this sensitive natural resource, only 503 non-commercial trips are allowed annually by the National Park Service (NPS). In 2015, over 20,000 applications were submitted. Far more people would love to raft the Colorado River than can be accommodated by the river without it incurring severe environmental harm. Given the trip's popularity record, it is entirely predictable that for the next year too, demand for permits will vastly outstrip supply. A permit to raft the Colorado River through the Grand Canyon is a scarce good.

In addition to cases of abundance and cases of scarcity, there exists a third type of case—bottleneck cases. Bottleneck cases obtain when it is not possible for each person to receive the good they demand immediately.¹¹ In bottleneck cases it would make everyone better off for the good to be allocated immediately, but various practicalities prevent this even though the good, over time, is abundant, thereby leading to a bottleneck. These practicalities include issues with the production and distribution of resources. Bottlenecks introduce the distributive significance of time. At t_0 the demand for the good is strictly greater than the supply for that good at t_0 . Yet at t_n , the demand for the good is lesser or equal to the supply. At time t_0 , not everyone who demands the good can get it because demand exceeds supply, but at time t_n the good will be available to everyone because supply is equal or greater than demand. Bottlenecks operate like scarcity at t_0 , and like abundance at t_n , with a gradual increase in supply in-between.¹²

¹⁰ We thank Leighton Reid for suggesting this example.

¹¹ To this extent, our definition of “bottlenecks” is broader than the other extant definition of bottlenecks in political philosophy: “the narrow places through which people must pass if they hope to reach a wide range of opportunities that fan out on the other side” (Fishkin 2014).

¹² There is also room for a fourth type of case where the change from scarcity to abundance happens instantaneously. We call these cases “manna cases” as there is a direct shift from scarcity to abundance, as with a sudden downpour of manna from heaven. Bottleneck cases are

Consider *Visiting Yosemite National Park*:

Each year, Yosemite National Park welcomes over four million visitors. Those that visit Yosemite Valley in summer can expect extremely high visitor concentrations, resulting in extended traffic delays. This is true even though, beginning May 2021, the NPS requires reservations to enter the park. Everyone who has a reservation for a particular day will be able to enter the park on that day, even if, as the NPS warns, they must wait an hour or more at entrance stations and up to two to three hours in Yosemite Valley itself. Entrance to Yosemite National Park, even if one has a reservation, is a bottleneck good.

The COVID-19 vaccine is a further example. In 2021 (t_0), when the vaccine was made available to the general population in the US, not everyone who wanted a vaccine could get it right away. At that moment, COVID vaccines were scarce in the US. However, given enough time the demand for the good was met, and anyone who wants to get a COVID vaccine in the US today can get one. COVID vaccines are now an abundant good in the US, and everyone who wants the good can have it at present (t_n). Focusing on the timeframe between abundance and scarcity, we recognize the case as one of a bottleneck. An essential feature of bottleneck cases is the temporal lag between demand and supply.

One may wonder whether bottlenecks are subject to a framing problem. From the point of view of an individual demanding a good, they may take themselves to be facing a case of scarcity rather than a bottleneck. For example, suppose an individual demanding a COVID vaccine is not aware of the fact that there will be enough vaccines to satisfy their demand for one. They may reasonably believe that they are confronting a case of scarcity, not a bottleneck. However, from the point of view of a decision-maker, they may be aware that there *will* be enough vaccines to supply everyone who needs one in the future. From their point of view, they are confronting a bottleneck case. As we demonstrate later on in the chapter, this flexibility in the framing of distributive scenarios is an advantage of the account. But for our

different from manna cases because there is a gradual shift between the states of scarcity and abundance. In the COVID-19 example below, between t_0 and t_n there is a gradual increase in the amounts of vaccine available at each time before everyone finally receives a vaccine. Moreover, there are cases in which the bottleneck never in fact turns into abundance, as in cases where there is a continuous queue that gets replenished with both demand and supply at some steady rate so that supply never meets demand.

purposes, we prioritize the point of view of the decision-maker, both because it is natural to focus on the epistemic perspective of those who will be distributing the good, but also for ease of expression.

What are these distinctions useful for? The distributive justice literature, and in particular the literature on the fairness of lotteries and queues, often operates under an assumption of scarcity, but sometimes conflates cases of scarcity and bottlenecks. This ultimately results in confusion. For example, when considering whether to use lotteries or queues to allocate goods fairly, Tyler John and Joseph Millum explain that they “mean this in two particular senses: [fairness] offers all candidates equal chances without regard to morally irrelevant characteristics, and [fairness] expressively signals the equal moral standing of all candidate recipients of a good” (John & Millum 2020). Yet both lotteries and queues can, at least ideally, meet such notions of fairness (Wasserman 1996). The question then is whether and when to use either allocation mechanism.

In Section 3, we demonstrate how cases of scarcity, rather than bottleneck cases, are usually assumed in the distributive justice literature, often only implicitly, and we then argue that lotteries are a more appropriate allocative mechanism in cases of scarcity. Similarly, in Section 4, we demonstrate how bottleneck cases are often implicitly used in the literature, and then argue that queues are a more appropriate allocative mechanism in bottleneck cases. As we demonstrate, cases of scarcity and bottleneck cases can each give rise to what we call secondary goods. These are goods that manifest when the demanded good is not available to everyone demanding it. We argue that it is appropriate to respond differently to the secondary goods in each case. The two secondary goods are “chances of receiving the good” and “lesser waiting time for receiving the good.” Lotteries are best placed to distribute chances in cases of scarcity, while queues are best placed to distribute waiting time in bottlenecks. We address some complications in Section 5 and conclude in Section 6.

3. Scarcity and Lotteries

Broadly, discussions of distributive justice tend to assume conditions of scarcity.¹³ The most popular method in the philosophical literature for

¹³ Rawls (1999), for example, refers to Humean idea of “conditions of moderate scarcity” when outlining the “circumstances of justice.” Otsuka and Voorhoeve (2009) consider cases of scarcity where it is not possible to satisfy the claims of each affected agent, when motivating their egalitarian position.

distributing scarce resources is the use of randomization procedures, such as lotteries.¹⁴ To our knowledge, every major account of the fairness of lotteries deals with cases of scarcity, rather than bottlenecks (Sher 1980; Broome 1984; Kornhauser & Sagar 1988; Broome 1990; Kamm 1993; Saunders 2008; Stone 2011).

In cases of scarcity, at least some unfairness is inevitable if we wish to maximize welfare.¹⁵ Demand necessarily exceeds supply, and so some will receive the good while others go without. Cases of scarcity make possible two types of considerations, what we term “primary” and “secondary.” What is of primary concern is that individuals who demand the good actually get the good. This is not possible to satisfy in cases of scarcity; there will necessarily be a conflict with respect to the primary consideration—that an individual actually receive the good. Nevertheless, in scarcity cases, it is possible to distribute something of value fairly, namely chances of receiving the good. This is a secondary consideration that can be invoked when it is not possible to satisfy the primary consideration. In a sense, this secondary consideration is always met even when all who can demand the good can actually receive it, since each will trivially have a chance of “1” of receiving the good. But in cases of scarcity, the default is that not everyone can receive the good and so at least one person’s chances of receiving the good will be lower than 1. Consequently, the value of chances is more clearly manifested in cases of scarcity.

Lotteries are seen as a paradigmatically fair procedure for cases of scarcity.¹⁶ Peter Stone, for example, states that a necessary circumstance for the use of a lottery in order to achieve allocative justice is that “there is not enough of the good to satisfy the claims of all these individuals with equally and maximally strong claims” (Stone 2011, 278). George Sher captures the general idea of what we can call the “lottery requirement”: “It is generally agreed that when two or more people have equal claims to a

¹⁴ There are a number of objections to the fairness of lotteries, however (Hooker 2005; Fumagalli 2022).

¹⁵ Some element of fairness can always be achieved by refraining from distributing any of the good at all. In such a case everyone is treated fairly, as King Solomon proposed to do when confronted with two women who claimed the same baby as their own (Kings 3:15–28). It is also an implication of John Broome’s theory of fairness as proportionate satisfaction of claims that in cases of scarcity it may often be fairer to withhold a good than give it directly to one of the candidates, as this would allow for a more proportionate treatment of claims (Piller 2017). Nevertheless, this increase in fairness comes at the expense of welfare, because withholding the good is welfare reducing.

¹⁶ Henning (2015) lists a number of authors who believe that lotteries are in fact *required* in scarcity cases where the good is one’s life being saved.

good that cannot be divided among them, the morally preferable way of allocating that good is through a tie-breaking device, or lottery, which is fair” (Sher 1980, 203). Again, there is reference to *a* good. This is the typical format of examples in the lottery literature: two individuals who each have a claim on *one* indivisible resource (Broome 1984; Broome 1990). Others still refer to the good *itself* as scarce; for example, Ben Saunders writes that he defends the justice of lotteries when distributing “non-divisible, *scarce goods*—such as school places, jobs or organs—between equal claimants” [italics added] (Saunders 2008, 359). Rather than deeming the good as scarce and viewing scarcity as an inherent attribute of the good itself, it is more appropriate to view the case itself as one of scarcity because the scarcity arises due to the particular relationship between supply and demand.

The existence of the secondary good of chances gives rise to the possibility of a “two stage” procedure whereby we check whether it is possible to distribute a good based on primary considerations, and if not, we can allocate the good via chances. In cases of scarcity, there will necessarily be a conflict of interest such that demand is not possible to be simultaneously satisfied with respect to primary considerations. However, cases of scarcity give rise to the possibility of the “secondary good” of chances.

We outline three reasons for thinking that lotteries, in general, are the most appropriate mechanism for allocating resources in cases of scarcity. First, a natural motivation for lotteries in cases of scarcity is that a chance of receiving a good can be distributed equally when the scarce good cannot. This position is called the “distributive view” of lotteries;¹⁷ when a proportionate allocation of the good between potential recipients is not possible, a lottery is able to divide what can be divided, namely the chance of receiving the good. So, lotteries are naturally adept at handling the secondary consideration of chances since they naturally trade in the distribution of chances. On this view, it is the secondary good of chances that is distributed.

Second, a chance of receiving a good is indexed to the good itself. It is something directly related to the good that can be offered as supplementary replacement. In this sense, the good of chances is “secondary” in that a chance of receiving some good is second-best to actually having that good. For example, suppose that Ann demands a medicine. It would be best for Ann to actually have the medicine. Now suppose that the pharmacist tells

¹⁷ Broome (1990: 97–8) also introduces a similar position through the idea that a lottery may provide “surrogate satisfaction” for a claim to a good. The distributive view more generally has come under recent criticism, including (Henning 2015).

her that they are currently out of stock, but it is possible to purchase a lottery ticket with a 90% chance of receiving the medicine later that day. It would be rational for Ann to value the chance of receiving the medicine and purchase the ticket (for at least some prices), even though she is not guaranteed to receive the medicine. Support for this general approach can be found in John Broome's account of the fairness of lotteries, where he refers to the idea that lotteries can provide some "surrogate satisfaction" for claims to a good (Broome 1990, 97–8). By entering a claim to a lottery, there can be a partial satisfaction of the claim in virtue of receiving a chance. This approach is open in cases of scarcity but not abundance. For our purposes, we do not need to settle on a particular account of the fairness of lotteries, but rather demonstrate that the allocation of chances is the most appropriate way to distribute the secondary consideration in cases of scarcity.

Thirdly, and aside from a consideration of the nature of chances themselves, we may think that a reasonable response to expected unfairness is to try to render such circumstances fair. For example, Broome writes that "If a good or bad cannot be distributed equally, it sometimes seems a good idea at least to distribute it randomly. Randomness appears to be a way of bringing some fairness into an inherently unfair situation" (Broome 1984, 40). When faced with scarcity, and the fact that not everyone who needs a particular good can receive it, we face a situation of unfairness. Chances, and their distribution can play a role in mitigating that unfairness (Diamond 1967; Broome 1990). For example, if there is only one medicine and two potential recipients, just giving the medicine outright to one of the recipients may be deemed unfair, compared to the alternative of giving each person an equal chance of receiving it. Equal chances facilitate equal treatment.

Another option for dealing with the unfairness that arises in cases of scarcity is to use queues to distribute the primary good. However, in cases of scarcity queuing has at least three significant drawbacks that affect the fairness and efficiency of the allocation.

First, the process of determining who is "first" in the queuing procedure can be arbitrary. For example, the "first person" in a queue for an ICU bed could mean either the first person to have fallen ill, the first person to arrive at the hospital, or the first person to be diagnosed (John & Millum 2020). As such, there is often an arbitrariness about who is "first" based on the criteria that is used to create a queue. This arbitrariness can be problematic when on different criteria there will be a different person who is "first." Each may therefore think that they are "first" and should therefore receive the good, if the allocators decide to use a queuing mechanism. If the allocator selects the

criterion of “first to have fallen ill,” then those who were first in the hospital and first to be diagnosed can complain that an alternative criterion ought to have been used, namely the criterion that would have placed them first in the queue.¹⁸

A lottery to decide which criterion is used to determine who is first will select the criterion in a way that does not appeal to arbitrary features. But the result of such a lottery is principally equivalent, and in some cases actually equivalent to deciding which person receives the kidney. For example, suppose that there is one kidney, three different criteria for determining who is first in line and three different people, each of whom would be first on one of the criteria. A lottery to determine which criterion should be used will be a surrogate for determining which particular individual will receive the kidney outright. In such cases, if a lottery is permissible for determining which criterion ought to be used, then a lottery ought to be permissible for determining who gets the kidney.

The second drawback is that queues are less effective than lotteries at screening out morally irrelevant differences that might influence distribution in cases of scarcity.¹⁹ To illustrate this point, consider the example of a refreshments table set out outside a meeting room. The allocators do not know how people will arrange themselves inside the room, and the allocators do not know in what order the people will file out to the table once the meeting is over. As such, the allocators have no reason to suspect that any particular person will be first in line. John and Millum argue that in such cases, queues tend towards perfect fairness (John & Millum 2020, 198). This is because it is epistemically equiprobable that each person will be first in line. There are no good reasons for thinking that any one person will be ahead of another.²⁰

In this case, epistemic equiprobability is established with reasoning via the “principle of insufficient reason.” Since there is no good reason to think that any particular person will be first out of the room, it is reasonable to think that everyone has the same chance. Although, of course, this will not in fact be true, because meeting rooms are configured in such a way that

¹⁸ One solution is to appeal to established conventional rules, however arbitrary they are. This would solve the problem in one regard as everyone will come to accept that “first to have fallen ill,” for example, will be the determinant of the queue.

¹⁹ See, for example, the “prophylactic view” outlined by Wasserman (1996). This points to a further argument in favor of lotteries, namely that they help screen out bias or partiality.

²⁰ It is important to note that this example implicitly assumes a bottleneck rather than actual scarcity because it is implicitly assumed that everyone in the queue will get their refreshments, eventually.

particular people are seated far away from the door and some are close to the door. John and Millum argue that such epistemic equiprobability is sufficient for the fairness of queuing mechanisms, but epistemic equiprobability arrived at in this way does not mean that there are in fact equal chances of being first. A lottery device on the other hand can assign precisely equal chances. Appealing to epistemic equiprobability through the principle of insufficient reason may also lead to information-aversion, where in order to preserve the fairness of the first come first serve mechanism, it would be best to not learn that some individuals are or may be placed earlier than others. In the refreshments table case, a clever and motivated audience member can easily “game the system” by sitting near the door or leaving early. While the conference organizers (the refreshment allocator) rightfully allocate epistemic equiprobability, morally irrelevant considerations can easily creep in.

Third, in cases of scarcity there are two options when it comes to queuing—either the individual queuing knows where the cut off is between those who will receive the good and those who will not, or they do not. If the individual contemplating queuing does not know where the cutoff is, they will not know, while in the queue, whether they will receive the good or not. They might be queuing, wasting precious time, only to end up with nothing. In effect, they will be waiting forever. If the individual contemplating queuing can predict whether their place in the queue guarantees that they will be allocated a scarce good and they know that they have not made the cutoff, this predictability entails that they should leave the queue, since they will not be allocated the good. However, when everyone who does not make the cutoff leaves the queue, this entails that everyone who is now demanding the good and remains in the queue is guaranteed the good. This *de facto* transforms the situation from one of scarcity to one of abundance (or a bottleneck), since the supply can meet this lesser demand.

As such, we can see that lotteries are a more appropriate mechanism than queues for distributing goods in scarcity cases. Firstly, the assignment of chances has a more plausible rationale than waiting time; secondly, lotteries are better at screening out morally irrelevant and arbitrary features of a distributive scenario when allocating a good, and thirdly, lotteries distribute chances more fairly than queues distribute waiting time when the good is scarce. When cases of scarcity arise because there is not enough of the good to meet demand, a lottery mechanism is more appropriate than a queuing mechanism.

4. Bottlenecks and Queues

As we argued in Section 2, distributive questions with respect to the good itself do not arise in cases of abundance. There is no fairness worry with respect to the primary consideration itself. However, bottleneck cases do give rise to the secondary consideration of waiting time due to the time lag between the demand for, and receipt of, the good that is the primary consideration. It is this secondary consideration that arises in bottleneck cases—that of shorter waiting times—that can be distributed more or less fairly.

Consider the case of a queue at the grocery store. What is at stake cannot be purchasing the groceries, since it is clear that barring some extremely unlikely situation those waiting in the queue will all check out eventually. Instead, what is at stake has to do with a secondary consideration that most of us have when joining the checkout line—checking out as quickly as possible. Longer waiting times are generally viewed as something bad, and shorter waiting times are generally viewed as something good. The distributive concern is with the secondary consideration of waiting time, not the primary consideration of checking out, and it only arises because we cannot all checkout immediately. It is a bottleneck case.

In this section, we argue that in bottleneck cases queues are a more appropriate distributive mechanism than lotteries because they are the fairest way to distribute the secondary consideration of waiting time efficiently.²¹ For perfect fairness we might aim to hoard the good until there is enough of it to supply all the demand simultaneously, in principle transforming the case from a bottleneck one to a manna case. If we are manufacturing a widget and the demand for it currently outstrips supply, which will eventually manage to catch up (think of something like the newest iPhone to come out), it is possible to simply delay supplying of the good to anyone until there is enough for everyone.

If the demand for the primary good begins simultaneously, then introducing a latency will result in a fairer distribution of the waiting time (the prospective distribution of the primary good is already fair, since it is abundant and all the demand will be met). The downside with introducing

²¹ Along similar lines, John and Millum (2020: 181) have recently provided what amounts to the strongest philosophical defense for queues in general (rather than just in bottleneck cases), making the case that queues are “relatively efficient, maximize[s] distribution equality relative to other Pareto efficient distributions, and treat[s] candidate recipients fairly.”

such a latency is that it is not an efficient way to promote welfare. While supplying the goods that are available as soon as possible to some of those demanding them reduces fairness, the tradeoff in terms of welfare usually gives us a sufficient reason to do so. Moreover, at best this solution only holds true when the demand all begins at the same time. If demand for the good is staggered, then introducing a latency so that everyone will receive the good at the same time entails that some individuals will wait longer from the moment their demand started than others, meaning that waiting time will intentionally be distributed unevenly.

Regardless of any consideration of fairness, for many bottleneck cases introducing an intentional latency is not practical, since many bottlenecks are not caused by production needing to catch up with demand, but rather by distribution challenges that make it impossible to simultaneously distribute of the good. Given that it takes an airplane passenger several seconds on average to reach for their carryon and start their move to debark the plane, it is simply impossible to avoid passengers having to endure some waiting before they get to leave the plane. Some passengers will get to leave sooner than others, and waiting time will need to be distributed. When distributing COVID vaccines, it takes time to administer a shot. When purchasing groceries in a supermarket, it takes time to check out one's groceries. This distributional challenge is the one that often gives rise to the bottleneck cases, and no amount of hoarding can correct for the fact that the distribution of the good to those demanding it will not be simultaneous.

Why not distribute goods in bottleneck cases by a lottery, just as in cases of scarcity? Why not think that the arguments in defense of lotteries in cases of scarcity generalize? The reason is that, unlike in scarcity cases, in bottleneck cases the likelihood that any given individual will receive the good is the same, and is equal to 1. In bottleneck cases there is ultimately no scarcity, and all the demand for the good will be met by the supply. A lottery when everyone is a winner is meaningless. The only way to make sense of applying a lottery in bottleneck cases is to conceive of it as a lottery for the secondary good of less waiting time rather than for the primary good. The problem with such lotteries is that they can introduce large disparities in how long people end up waiting for the good. If a person arrives into the lottery pool early and is continuously unlucky, they might wait significantly longer than a person who just showed up and enjoyed good luck. A lottery system may also be unfair for bottleneck cases, for example, a lottery for season tickets for a popular football team. If everyone that waits for such a ticket will eventually receive one, even if it takes years, holding periodic lotteries will

likely mean that some will “win” the lottery before those who have waited longer receive the ticket. In effect, the winners will not need to wait at all, while others wait, unluckily, for decades.²²

This unfairness, however, is avoided when using queues to allocate and distribute waiting times. The main reason to think that queues are a fairer distributional mechanism than lotteries in bottleneck cases is, as John and Millum also convincingly argue, that while lotteries can approach queues in fairly distributing waiting times, queues can do so reliably and consistently in some specific cases. As they make clear:

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 waiting for a resource is (cardinally) an equal counterfactual harm for each person compared to receiving the resource. (John & Millum 2020, 195–6)

The standard picture of a grocery store queue, passport control queue, or a breadline, all involve queues that are continuous, with people joining the queue all the while the goods are being allocated. For such bottlenecks it is possible to fairly distribute waiting times. In such queues, those near the beginning of the queue have waited for the good for a while already. Conducting a lottery to determine who will receive the good next in bottleneck cases in which people already have different waiting times is unfair. When demand for the good trickles in at some rate, R_n , and supply of the good trickles in at some rate, R_m , queues will be the most efficient way to allocate waiting time in a way that is equal at the limit. The fairest way to allocate the secondary consideration of low waiting times is to attempt to equalize it. When $n = m$ and those joining the queue are doing so at a rate equal to the rate at which goods become available, then everyone can wait an equal time. When $n \neq m$ and the rate at which people join the queue is different from the rate at which goods become available, a queue will be the best means to approach equal waiting times. If R_n or R_m are non-uniform, and sometimes the queue moves faster than other times, ordinality of waiting time is maintained, although a queue cannot guarantee that waiting times themselves will be as close to equal as possible.

²² We thank Carl Knight for proposing this example.

It is true that, as John and Millum argue, waiting time does not have intrinsic moral significance, and “the fact that someone has waited longer in a queue for a scarce good is not intrinsically morally significant” (John & Millum 2020, 180). Waiting time, in and of itself, does not matter intrinsically. Moreover, sometimes waiting time does not negatively affect the individual, and so, other things equal, an individual with more waiting time is no more burdened than an individual with less. There are even cases in which waiting time can seem to have a positive effect on the individual. For instance, waiting for a gift or a pleasurable experience. Queuing can also have beneficial elements, like the feeling of solidarity one might experience with those waiting with them in the very long queues to vote in some districts in the US.

Nevertheless, such cases are the exceptions that show the rule. Overwhelmingly, waiting is considered a bad. John and Millum concede, for example, that waiting time “very regularly correlates with something that is intrinsically morally significant: unpleasant experience” (2020, 181). This correlation between waiting time and unpleasant experience is regular enough that, from the perspective of the allocator, waiting time is a harm they ought to aim to distribute fairly.²³ It is generally true that long waiting times are perceived as a cost and a burden, even if they are so only indirectly.

Moreover, we can see that it is the temporal aspect of the time waited, rather than the ordinal position in a queue that is really what we care about, absent considerations of custom and etiquette. When searching for the shortest queue, in the sense of it having the least amount of people in it, what matters is not to have the lowest ordinal number of people ahead of us. What matters to us is that we join a queue with fewer people in it because it is a reasonable proxy for a queue with less waiting time. Since we know we will checkout at the supermarket at some point, the relevant secondary consideration is to minimize our waiting time. It is true that we would not like to be skipped ahead in line. But this has to do with a sense that established rules or norms have been broken. To explore this intuition, Consider *Queue Switchers*:

You are standing in a queue in the grocery store behind a person with a cart full of groceries. It would take a full five minutes for them to get checked

²³ We do tend to care more about waiting time for people for whom a longer wait means outright suffering (e.g. waiting for an organ transplant) than we care about waiting time for people who are bored at a queue (e.g. waiting for a ride in the amusement park), but these are merely differences in degree, not differences in kind.

out. Suddenly, they decide to leave the queue, not before telling two other people with only one item each, that they can take their place in the queue. As a result, you wait a full three minutes less.

Of course, there are a variety of social norms that are violated in this scenario, and we might indeed be upset that this was done unilaterally. Perhaps switching your place in the queue with one other person is socially acceptable, but surely it is not acceptable to switch one's place with two unrelated people. At the very least they should ask those who are standing behind them if they agree. But ultimately, what matters in a grocery store checkout line is to conclude the transaction as quickly as possible. It would be reasonable to choose the queue with ten people, each with two items, over the queue with two people, each with a huge cart full of groceries and a handful of coupons, because we would expect to wait less in the ordinarily longer queue. That we would feel wronged by someone cutting ahead in line, even if it shortens our wait somehow does not invalidate this. It simply demonstrates that once a norm is established, we find those who violate it injurious.

Queues are superior to lotteries in bottleneck cases. In bottleneck cases there is no inequality in the distribution of the primary good itself, since everyone will get it. The potential inequality only arises with respect to waiting time, not the good itself. It is waiting time that must be distributed fairly. Consequently, considerations raised in this section only apply to bottleneck cases.

5. Complications

Several objections might be raised. First, it could be pointed out that while for many bottlenecks people join the pool of those demanding the good continuously, as would be the case with cars on a highway that narrows down from three lanes to two, there are plenty of cases in which there is no temporal primacy among those who join the pool at different times.

Consider again the case of the COVID-19 vaccine. Demand for a vaccine arose at least as soon as they came into existence. The supply of such a vaccine in the US only began around December 2020. While the COVID-19 vaccine supply came trickling in, we can assume that demand for the vaccine was near universal, i.e., that the majority of adults around the world (and specifically in the US) were interested in obtaining a vaccine. In effect, the

whole world joined the queue for the COVID-19 vaccine at the same time, when COVID-19 became a pandemic. Even if the variety of morally relevant considerations such as age, profession, and medical condition are prioritized, in the US alone there are at least a hundred million adults whose demand for the vaccine pragmatically started at the same time. Since, even once the vaccine opened up to all adults in the US on March 19, 2021, there was insufficient supply to meet demand, there is no reason to think that whoever entered the queue did so any earlier than anyone else. But some waiting time is required.

Our response is that lotteries could be used in a different way than how they are used in scarcity cases. Instead of using a lottery to determine *who* will get the good, in bottleneck cases we could have a lottery to determine *when* one will get the good. The lottery is not directly for the in-demand good, which is the primary consideration as in cases of scarcity, but for shorter waiting times, which is a secondary consideration. Those who “win” the lottery will wait less for the good than those who “lose” the lottery. In cases like that involving the COVID-19 vaccine, waiting time cannot be distributed *equally* at all. But waiting time can nonetheless be distributed *fairly*.

Consider what such bottleneck cases amount to; there is a good that everyone will get, so the good itself is not scarce. But the good cannot be distributed instantaneously. However, there is no morally relevant difference on the basis of which to organize the queue. Some people will need to wait longer than others, with no good principled way to determine who does so. Thus, low waiting time itself is a secondary, scarce, good. Only some individuals will enjoy low waiting times while others will suffer long waiting times. As we discussed in Section 3, the appropriate way to distribute a good in cases of scarcity when there are no morally relevant considerations is through a lottery. But the lottery in this case is for low waiting times, not the good itself. Waiting time is a secondary consideration. It just now adds a third consideration—the fair distribution of the secondary consideration.

In cases in which the demand begins at an instant, such as demand for COVID-19 vaccines, people queuing up for refreshments after a conference talk, the doors open at a store on Black Friday, small and morally irrelevant differences in proximity, internet savvy, or being able to run faster than someone else, can dictate the ordinality of the queue and consequently potentially large differences in waiting time. For such cases lotteries for the secondary good of low waiting time are appropriate.

Another complication occurs when bottlenecks arise in cases of scarcity. Suppose, for example, that there are twenty people who each demand a good. There are only ten goods, and only one good is released per hour. In this case, not everyone will receive a good and not everyone will receive it at the same time. Furthermore, a bottleneck appears to occur concurrently with scarcity. One might think that our proposed framework faces an impasse in such mixed cases. However, we can simply recognize two relevant conceptual stages. First, there are conditions of scarcity. There will only be ten goods and there are twenty people. Given this scarcity, it is appropriate to use a lottery to determine who will receive a good. This will solve the problem of scarcity. For the subset of ten individuals who will get the good, the case is no longer one of scarcity. The scenario can now be treated as a bottleneck case with ten individuals and ten goods. If the demand for the good is staggered, the queue should form in an ordinal manner. If the demand for the good occurs simultaneously, then, as we have argued previously, another lottery is appropriate to determine the order of the queue.

A third potential objection is that we seem to assume a unity we are not entitled to. We treat “good at t_n ” as the same as “good at t_m .” This can be thought as lacking in the rigor that Mas-Colell et al. require when they write that “time (or, for that matter, location) can be built into the definition of a commodity. Rigorously, bread today and tomorrow should be viewed as distinct commodities” (Mas-Colell et al. 1995, 18). After all, as Fisher points out, usually most of us are not indifferent between consuming the same good now or in the future, and we display a preference for present over future goods (Fisher 1930). However, we *sometimes* assume such a unity for good reason. Rather than assume time separability of goods and frame the consumption of goods at different times as distinct goods, we wish to sometimes separate the consumption of the goods from the time spent waiting for them. Separating the consumption of the good and waiting for that good is particularly important in the distributional justice context because questions become moot if we cannot treat “good at t_n ” and “good at t_m ” as commensurable for distributional purposes.²⁴

²⁴ This is not to say that it is *always* appropriate to reject time separability of goods. A good case in point has to do with clothing and fashion. It would be unwise to conceptualize skinny jeans or bell bottoms separate from the time they can be distributed and consumed given how much their consumption is tied up with the time at which they are considered fashionable. We thank Samuel Mortimer for this comment.

Suppose a parent can distribute a non-divisible chocolate bar to their kids every ten minutes (they need to go to the refrigerator each time and can only carry one chocolate bar at a time). The parent holds a lottery at t_1 and then distributes a chocolate bar to one child at t_1 . The parent then goes to get another chocolate bar. When the parent comes back, the first child already consumed their chocolate bar and demands another. When the parent says that it is only fair that the second child should get it, the first child protests that what was distributed in the past bears no relevance to the distribution of this chocolate bar, because they are conceptually different. Chocolate bar at t_1 is not the same good as chocolate bar at t_2 , and a lottery must be had for this new and distinct good. If we accept the temporal separability of goods, there is some merit to the child's claim.²⁵

One could run a parallel argument to this chocolates case, but instead focus on something more morally significant, for example a partial tuition subsidy for college. Two students (in similar circumstance) vie for a partial tuition subsidy, and after student₁ gets a subsidy at t_1 and student₂ worked overtime to make up the difference, student₁ claims at t_2 that past distribution bear no relevance to the present, and a lottery is again appropriate.

When we accept the temporal separability of goods by indexing goods to a time, we lose the distributional depth that raises the distributional problem. If goods are always indexed to time, then we lose something that matters for fairness. Something is lost if we carve up the distributional problem in a way that prevents us from addressing the bottleneck, because at every given moment it is, strictly speaking, either a case of scarcity or a case of abundance. If bottleneck cases are dismissed as solely cases of scarcity at t_0 and as cases of abundance at t_n , then one is not sensitive to the temporal element of the distribution problem. We only see the scenario as a bottleneck case if we view it over some period of time.

Even Mas-Colell et al. concede that viewing goods as purely instantaneous is neither practical nor appropriate:

Although commodities consumed at different times should be viewed rigorously as distinct commodities, in practice, economic models often involve

²⁵ If the reader finds it difficult to imagine such a scenario, the reader must have little experience with children and their ability to sophistically advocate for themselves when sweets are involved.

some “time aggregation.” Thus, one commodity might be “bread consumed in the month of February,” even though, in principle, bread consumed at each instant in February should be distinguished.

(Mas-Colell et al. 1995, 18)

If we view the timeframe over which we are concerned narrowly enough, we can always frame away the bottleneck, either in favor of viewing it as a case of scarcity or as a case of abundance. While it is, of course, possible to do so, it is not helpful to do so. Framing the situation narrowly enough allows us to avoid treating the case as a bottleneck case, but it also causes us to ignore the relevance of the broader picture that a larger timeframe provides. This has ramifications for how we conceive of fair distribution.

6. Conclusion

In this chapter we argued that what distributive mechanism is appropriate will depend on the specification of the case that the distributing agent faces. This became apparent once we introduced a new element of distributive concern: bottleneck cases. These cases feature as an alternative alongside scarcity and abundance. In cases of abundance, fair distribution is not a question; in cases of scarcity, the use of a lottery is most appropriate; in bottleneck cases, the use of queues is most appropriate. The appropriateness of each mechanism is down to their fairness and efficiency in each distributive predicament. This chapter’s contribution is in making salient the distinct value of recognizing the temporal nature of distributive scenarios. Our account makes clear when we should use queues and when we should use lotteries to allocate goods.²⁶ It makes clear why sometimes we think a queue-based system is unfair and why we sometimes think that lotteries are inappropriate. This framework has potentially fruitful applications to cases of high moral stakes, such as medical resource allocation (e.g. kidney allocation), immigration policy, and housing vouchers, where the appropriate allocation mechanism will depend on to what extent there is scarcity or a bottleneck.

²⁶ For an example of a specific practical application of some of these ideas see Hersch (2022), in which he argues for a Random Selection for Service (RSS) mechanism over a First In First Out (FIFO) mechanism for financial exchanges to match standing limit order with incoming market orders.

This framework can be used to defend some uses of queues in societal contexts in which they are already established, for example for queues for buses when everyone will eventually get a seat or queues for medical treatment when everyone will eventually be seen. The arguments we offer can also be used to defend some current uses of lotteries in societal contexts. For example, lotteries are sometimes used in school admissions, such as charter schools in the US and in some districts in the UK.²⁷ Given that whole school year cohorts will begin their schooling at the same time, it is a case where individuals simultaneously confront the scarcity of school spots (assuming that for particular schools there will be more students that apply than places). This is a case of scarcity. As such, a lottery is appropriate to remedy the scarcity.²⁸

Our framework also has implications for cases where lotteries are used but queues would be more appropriate. Consider the case of airport passport control. Some queueing systems provide multiple sub-queues, one for each passport control booth. This creates a quasi-lottery because individuals need to select which sub-queue of the many available to join. From the point of view of the individual, they do not know which sub-queue will get them through passport control fastest, so deciding which sub-queue to join is a *de facto* lottery. Some sub-queues move quickly and without a hitch, while others might have a person without the proper visa and so take an inordinate amount of time. Often it is the case that some who arrive later than others will be served before those who arrived earlier. Our framework suggests that because individuals will typically arrive at the queue at different times, it is fairest to use a single queueing system rather than such a multi-sub-queue system.²⁹

Alternatively, our framework helps single out cases where a queue is used but a lottery would be more appropriate. One example has to do with a daycare that also runs a “camp” during the regular school breaks.³⁰

²⁷ A lottery was introduced in the English city of Brighton to allocate oversubscribed school places in 2007, “War Over School Boundaries Divides Brighton,” 01/03/07, *The Guardian*: <https://www.theguardian.com/uk/2007/mar/01/schooladmissions.topstories3>. Stone (2013) has argued that in favor of a lottery to break ties between potential university students, and Sandel (2020) has recently proposed a lottery for the allocation of university places to students.

²⁸ One could use queues and admit students on a first come first serve basis, closing registration once the first one hundred students registered. However, as we argue in this chapter, such a system would be less fair than a lottery-based system.

²⁹ The US immigration is usually organized in such a quasi-lottery system, whereas UK immigration has a single queue that is broken up as people arrive at the beginning of the queue.

³⁰ This example is based on the personal experience of one of the authors, who actually benefits from the current queue-based system since they usually work from their computer, yet still recognize the system as less fair.

While the dates at which the breaks are scheduled are known well in advance, the daycare administration sends out a registration notification email for the camp at some a few weeks before the camp begins. Parents do not know when to expect this email. Registration is on a first come first served queuing basis from the moment the email is sent out, and spots fill out very quickly. Those that work by computers are much more likely to quickly respond and get their kid signed up for camp. Many parents at the daycare find this system inappropriate. Our account makes clear why their judgment makes sense. Spots at the camp are scarce and a queue-based system does not suit cases of scarcity. The daycare should change its allocation system.³¹

Bottleneck cases have unique attributes and lumping them in with either cases of scarcity or abundance has resulted in confused intuitions regarding fairness in distribution. Recognizing that when it comes to distributive justice the temporal aspect matters, and that between abundance and scarcity we can find the unique context of bottlenecks helps make sense of how we think, and ought to think, about fairness in distribution.³²

References

- Brake, Elizabeth (2020), "Price gouging and the duty of easy rescue," *Economics and Philosophy* 37: 1–24.
- Broome, John (1984), "Selecting people randomly," *Ethics* 95: 38–55.
- Broome, John (1990), "Fairness," *Proceedings of the Aristotelian Society* 91: 87–101.
- Carens, Joseph (1987), "Aliens and citizens: The case for open borders," *Review of Politics* 49: 251–273.
- Diamond, Peter (1967), "Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: comment," *Journal of Political Economy* 75: 309–321.
- European Parliament (2013), *Resolution on EU Citizenship for Sale*, available at: <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//TEXT+MOTION+P7-RC-2014-0015+0+DOC+XML+V0//EN>.

³¹ In fact, while this article was under review, the daycare did change its system to a lottery-based one, unprompted by the authors.

³² We would like to thank audiences at the UPenn Zicklin Center Normative Business Ethics Workshop 2022, the 10th Annual Workshop for Oxford Studies in Political Philosophy, the American Philosophical Association Eastern Division meeting in Montreal 2023, and to members of the Virginia Tech Ethics/Social & Political (ESP) Research Cluster Group.

- Fisher, Irvine (1930), *The Theory of Interest*, New York: The Macmillan Company.
- Fishkin, Joseph (2014), *Bottlenecks: A New Theory of Equal Opportunity*, Oxford: Oxford University Press.
- Fumagalli, Roberto (2022), "We should not use randomization procedures to allocate scarce life-saving resources," *Public Health Ethics* 15: 87–103.
- Goodin, Robert (2001), "Managing scarcity: toward a more political theory of justice," *Philosophical Issues* 11: 202–228.
- Henning, Tim (2015), "From choice to chance? Saving people, fairness, and lotteries," *Philosophical Review* 124: 169–206.
- Hersch, Gil (2023), "Procedural fairness in exchange matching systems," *Journal of Business Ethics* 188: 367–377.
- Hidalgo, Javier (2016), "Selling citizenship: A defense," *Journal of Applied Philosophy* 33: 223–239.
- Hooker, Brad (2005), "Fairness," *Ethical Theory and Moral Practice* 8: 329–352.
- Hume, David (1777), *An Enquiry Concerning the Principles of Morals*, Section III. Part I.
- John, Tyler and Millum, Joseph (2020), "First come first served," *Ethics* 130: 179–207.
- Kamm, Frances (1993), *Morality, Mortality Vol. 1: Death and Whom to Save from It*, Oxford: Oxford University Press.
- Knight, Carl (2011), *Responsibility and Distributive Justice*, Oxford: Oxford University Press.
- Kornhauser, Lewis and Sagar, Lawrence (1988), "Just lotteries," *Rationality and Society* 27: 483–516.
- Lamont, Julian (1997), "Incentive income, deserved income, and economic rents," *Journal of Political Philosophy* 5: 26–46.
- Lamont, Julian and Christi Favor (2009), "Price gouging in disaster zones: an ethical framework," *Social Alternatives* 28: 49–54.
- Mas-Colell, Andreu, Whinston, Michael, and Green, Jerry (1995), *Microeconomic Theory*, New York: Oxford University Press.
- Miller, David (1989), *Market, State, and Community*, Oxford: Clarendon Press.
- Milne, Heather (1986), "Desert, effort and equality," *Journal of Applied Philosophy* 3: 235–243.
- Otsuka, Michael and Voorhoeve, Alex (2009), "Why it matters that some are worse off than others: An argument against the Priority View," *Philosophy & Public Affairs* 37: 171–199.

- Piller, Christian (2017), "Treating Broome fairly," *Utilitas* 29: 214–238.
- Raiklin, Ernest and Uyar, Bülent (1996), "On the relativity of the concepts of needs, wants, scarcity and opportunity cost", *International Journal of Social Economics* 23: 49–56.
- Rawls, John (1999), *A Theory of Justice: Revised Edition*, Cambridge, Mass.: Harvard Belknap Press.
- Reding, Viviane (2014), "Citizenship must not be for sale", European Commission Speech 14/18, available at: http://europa.eu/rapid/press-release_SPEECH-14-18_en.htm.
- Robbins, Lionel (1935), *An Essay on the Nature and Significance of Economic Science*, London: Macmillan & Co.
- Sandel, Michael (2012), *What Money Can't Buy: The Moral Limits of Markets*, New York: Farrar, Straus and Giroux.
- Sandel, Michael (2020), *The Tyranny of Merit*, London: Allen Lane.
- Satz, Debra (2012), *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*, Oxford: Oxford University Press.
- Saunders, Ben (2008), "The equality of lotteries," *Philosophy* 83: 359–372.
- Shachar, Ayelet (2009), *The Birthright Lottery: Citizenship and Global Inequality*, Cambridge, MA: Harvard University Press.
- Shachar, Ayelet and Hirschl, Ran (2014a), "On citizenship, states, and markets," *Journal of Political Philosophy* 22: 231–257.
- Shachar, Ayelet and Hirschl, Ran (2014b), *Should Citizenship Be for Sale?* Florence: European University Institute.
- Sher, George (1980), "What makes a lottery fair?," *Nous* 14: 203–216.
- Snyder, Jeremy (2009), "What's the matter with price gouging?," *Business Ethics Quarterly* 19: 275–293.
- Stone, Peter (2011), *Luck of the Draw: The Role of Lotteries in Decision Making*, Oxford: Oxford University Press.
- Stone, Peter (2013), "Access to higher education by the luck of the draw," *Comparative Education Review* 57: 577–599.
- Van der Vossen, Bas and Jason Brennan (2018), *In Defense of Openness*, New York: Oxford University Press.
- Wasserman, David (1996), "Let them eat chances: Probability and distributive justice," *Economics and Philosophy* 12: 29–49.