



UvA-DARE (Digital Academic Repository)

Experimental studies on the psychology of property rights

El Haji, A.

Publication date

2017

Document Version

Other version

License

Other

[Link to publication](#)

Citation for published version (APA):

El Haji, A. (2017). *Experimental studies on the psychology of property rights*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

5. Trading Places

5.1. Introduction

It is well-known that queues where customers are served on a first-come first-served basis are an inefficient way to ration scarce service time. The reason is that the queuing order does not guarantee that customers with high waiting costs are served before those with low waiting costs. Extant empirical research that explores ways to improve the customer's experience of waiting in a queue focuses primarily on policies that shorten overall waiting time or improve the waiting experience (*e.g.*, Kumar et al. 1997; Pruyn and Smidts 1998). Policies that allow customers with high waiting costs to get ahead for a price might improve the experience for customers who can afford this service but harm the other customers in the queue (Zhou and Soman 2008). In this chapter we explore the possibility of allowing customers to, literally, trade places to increase a queue's efficiency. This would allow customers to get ahead in a queue while compensating other customers for their longer waiting times.

We study the efficiency-enhancing properties of two intuitively appealing auction mechanisms that facilitate customers' trading places using a laboratory experiment. Kleinrock (1967) shows that a queue's efficiency may be restored if customers' positions depend on how much they pay the server. In Kleinrock's model it is assumed that upon arrival each customer pays a 'bribe' to the server. The server places the customer in the queue behind all customers who offered a higher bribe and in front of those who paid a smaller bribe. Assuming that in the steady state customers

use the same bribing function that is strictly increasing in marginal waiting costs, Kleinrock shows that steady-state waiting times are minimized, which results in efficient queues.⁵ However, the assumption that customers use the same, strictly increasing bribing function may be a strong one in practice. It may be unrealistic if customers are heterogeneous in dimensions other than only marginal waiting costs (such as risk attitude or beliefs about others' waiting costs). Moreover, results from economic experiments show that, in general, human bidding behavior is 'noisy' so that even in a setting that satisfies Kleinrock's assumptions, inefficiencies are still likely to occur.⁶

In this chapter, we compare two auction mechanisms that could be used to determine the sequence of service to queued customers: the server-initiated auction and the customer-initiated auction. In the server-initiated auction, the server, when idle, invites each queuing customer to submit a bid. The server will then serve the customer who has submitted the highest bid. This customer shares her bid equally among each of the remaining customers in the queue. In the customer-initiated auction, a new arrival can sequentially trade places with customers currently in the queue. The arriving

⁵ Lui (1985), Glazer and Hassin (1986), and Afèche and Mendelson (2004) back up Kleinrock's (1967) result by showing that an efficient queue order emerges in a Bayesian-Nash equilibrium in settings where customers incur waiting costs that are linear in waiting time. Hassin (1995) shows this can be achieved with exponential waiting cost functions. Kittsteiner and Moldovanu (2005) generalize the equilibrium analysis, allowing for convex and concave waiting cost functions. See Hassin and Haviv's (2003) book for a discussion of some of this literature.

⁶ See Kagel (1995) and Kagel and Levin (forthcoming) for overviews of results from the experimental auctions literature. In most auction formats, inefficiencies arise because participants employ different bidding strategies, even after ample learning opportunities.

customer offers money to the queuing customers, from the back to the front. The current customers indicate simultaneously the minimum amount they are willing to accept. A new arrival trades places with a customer in front of her if and only if the latter is willing to accept her offer. This process stops as soon as the new arrival does not trade places with the customer in front of her.

We focus on the two particular auction mechanisms because both have intuitively appealing properties. First of all, an efficient ordering is feasible for both mechanisms if customers act non-strategically.⁷ The server-initiated [customer-initiated] auction implements the selection sort [insertion sort] algorithm that ensures an efficient queuing order if customers consistently submit bids equal to their marginal waiting costs. Moreover, both mechanisms can be straightforwardly used in a dynamic setting where customers arrive while the server is busy. In addition, both mechanisms are budget-balanced from the viewpoint of the customers, in contrast to Kleinrock's (1967) 'bribing mechanism.' As both auction mechanisms have the potential to decrease total waiting costs, they increase the 'pie' compared to a setting where customers cannot trade places. Because all gains-from-trade remain in the customers' hands, entry into the queue is not discouraged, in contrast to a mechanism where customers pay the server to obtain priority.⁸ Furthermore, as discussed below, the two auction

⁷ A priority queue is an example of a mechanism that cannot guarantee an efficient ordering. While opening a priority queue may improve the efficiency compared to the situation where only the original queue exists, inefficiencies still remain because the two queues may still be ordered inefficiently.

⁸ Yang et al. (forthcoming) study mechanisms where queued customers compensate an intermediary for the opportunity to trade positions.

mechanisms are predicted to differ in terms of attractive properties like efficiency and fairness. Finally, comparing the two mechanisms may reveal which mechanism is more attractive for marketing purposes in the sense that a firm offering relatively efficient or relatively fair queues may be more attractive for new potential consumers.

We compare the behavioral properties of the two mechanisms in a laboratory experiment. In contrast to Kleinrock (1967) and most of the theoretical queuing literature, we analyze the mechanisms in a static environment. In this environment there is a fixed and commonly known number of customers waiting in line to be served by the server. The server only opens as soon as all the customers have arrived in the queue.⁹ We have chosen this setup for two reasons. First, it is hard, if not impossible, to find analytical results for dynamic processes in a transient state, so that our experimental study would become a fishing expedition without clear testable hypotheses. Second, it is practically impossible to invite so many participants in a laboratory setting to implement a dynamic process that evolves reasonably close to a steady state.

We evaluate the two auction mechanisms along two dimensions: efficiency and perceived fairness. To develop testable hypotheses regarding efficiency, we derive the theoretical properties of the mechanisms in an independent private waiting costs model. In our model, customers face constant marginal waiting costs per unit of

⁹ As a consequence, our setting translates into a scheduling problem. Mitra (2001), Wellman et al. (2001), Feng (2008), Kayı and Ramaekers (2010), and Gershkov and Schweinzer (2010) also study auctions mechanisms used for job scheduling.

time. A customer's initial position is independent of her marginal waiting costs. We show that the server-initiated auction has an efficient (Bayesian-Nash) equilibrium, in contrast to the customer-initiated auction. The latter finding is not surprising in light of Myerson and Satterthwaite's (1983) impossibility result which shows that in a large range of settings efficient trade between an incompletely informed buyer and seller is not feasible. In our setting, customers at the front of the initial queue 'own' their position so that trade with late arrivals will not occur as often as efficiency requires. In contrast, users may perceive the customer-initiated auction as a fairer mechanism than the server-initiated auction because only the former grants them ownership rights over their initial position.

To examine the behavioral properties of the two auction mechanisms, we use two novel experimental protocols. Our first protocol implements induced waiting costs. Before bidding in the auctions, participants are privately informed of their own marginal waiting costs. Depending on the number of turns participants have to wait before being served, we subtract the resulting waiting costs from their starting capital. The efficiency gain resulting from the auctions can be readily measured because the induced waiting costs are known to the experimenter. The second protocol involves actual waiting. We used this protocol to determine the order by which participants could leave the laboratory. Participants vote for either of the two auction mechanisms and a majority rule determines which auction is actually implemented. In addition, participants were asked in a questionnaire to rate the auctions in terms of fairness on a seven-point Likert scale.

Besides studying the outcomes of the auction mechanisms, we also check whether psychological biases like endowment and sunk-cost effects have an impact on bidding behavior. We do so by varying the arrival process as part of our experimental design. On the basis of the literature, we conjecture that the endowment effect and the sunk-cost effect can simultaneously affect behavior in a setting where customers can trade places in a queue. The endowment effect occurs when the sheer possession of an object increases a person's value for it. Indeed, significant endowment effects (measured by a willingness-to-accept/willingness-to-pay gap) are observed in many other contexts.¹⁰ Anecdotal evidence suggests that people standing in line feel entitled to their queue position, which in turn could result in an endowment effect.¹¹ Specifically, if customers feel that they own their current position in the queue, they may be willing to bid a higher amount when their position is up for auction than standard theory predicts.

Someone falls prey to the sunk-cost bias if her decision depends on unrecoverable costs that are economically irrelevant for the

¹⁰ Knetsch (1989) and Kahneman et al. (1990) provide early examples.

¹¹ Mann (1969) observes queue jumping being discouraged in waiting lines for tickets to watch the "world series" of Australian rules football in Melbourne, Australia. Helweg-Larsen and LoMonaco (2008) find similar responses in a survey among fans of the Irish rock band U2 queuing for concert tickets. Milgram et al. (1986) let confederates impose themselves into queues in train stations and other public locations in New York and report customers' defensive reactions varying from expressing verbal objections to physical actions against the intruders. Oberholzer-Gee (2006) finds many customers willing to let someone jump the queue when offered a monetary compensation. However, when approached for a second time, all "individuals rejected my request, most of them appeared upset, some angry, a few outright hostile, suggesting that it was probably not safe to continue the experiment."

decision at stake.¹² Time spent waiting in a queue is such a sunk cost. Standard economic theory assumes that waiting costs do not affect a customer's willingness-to-pay for queue positions. In the case of a sunk-cost effect, a customer's valuation of queue position depends on how much time she has spent waiting in the queue. The existence of endowment and sunk-cost effects in a queuing setting implies that auctions that allow trading places cannot guarantee that the final queuing order is efficient.

Our main results are the following. First of all, the two auction mechanisms considered do not differ in a statistical meaningful way with respect to the average efficiency gain, irrespective of the arrival protocol. This is surprising in light of our theoretical findings that the server-initiated auction has an efficient equilibrium while the customer-initiated auction does not. In a deeper examination of our data, we do observe differences between the auctions in terms of efficiency gains: Efficiency gains are significantly greater [lower] in the server-initiated auction than in the customer-initiated auction if the initial queuing order is relatively inefficient [efficient]. Neither auction comes close to always reaching an efficient outcome. For the server-initiated auction, this result is rooted in noisy individual bidding behavior that is partly explained by a sunk-cost effect but not by a noticeable endowment effect. Noisy behavior in the server-initiated auction explains why efficiency gains are low and often even negative if the initial queuing order is already relatively efficient. In the

¹² See the seminal paper by Arkes and Blumer (1985). Phillips et al. (1991), Offerman and Potters (2006), Friedman et al. (2007), and Baliga and Ely (2011) present further experimental evidence on the sunk-cost bias.

customer-initiated auction, the queuing order remains relatively inefficient because customers bid more aggressively for their current position than arriving bidders do. In addition, we find evidence of a sunk-cost effect in the customer-initiated auction, which contributes to the auction mechanism's modest efficiency gain. On the positive side, the observed bidding behavior implies that it is unlikely that a customer will trade places if the queue is already in an efficient order. This explains why the customer-initiated auction outperforms the server-initiated auction if the initial queue's order is relatively efficient. Finally, when given the choice between the two auction mechanisms, participants tended to favor the server-initiated auction. This may be partly explained by participants evaluating the server-initiated auction as fairer than the customer-initiated auction.

This chapter speaks to several literatures. First of all, this chapter contributes to the marketing literature in that it studies customers' waiting experience. As far as we know, we are the first to experimentally study priority auctions in a queuing setting. Second, we add to the behavioral operations literature.¹³ Several papers within this literature examine queuing processes in the lab. Rapoport *et al.* (2004), Seale *et al.* (2005), and Stein *et al.* (2007) study participants' decisions as to when to enter a queue, if at all, to test whether participants' arrival times are consistent with Nash equilibrium predictions. Kremer and Debo (2012) examine queue herding in a setting where participants can enter a queue to obtain a good of an uncertain quality. We also contribute to the behavioral

¹³ See Bendoly *et al.* (2010) for a recent overview of this literature.

economics literature by examining the endowment effect and the sunk-cost effect in a setting involving queues.

5.2. Theoretical framework

Consider a queuing system where $N \geq 2$ risk-neutral customers, labeled $i = 1, \dots, N$, arrive sequentially in a queue to get served by a server. Each customer is privately informed of her waiting costs per unit of time, which we will denote by c_i . We assume that the c_i 's are independently drawn from a differentiable distribution function F on an interval $[\underline{c}, \bar{c}]$, $\bar{c} > \underline{c} \geq 0$, with $F'(c) > 0$ for all $c \in [\underline{c}, \bar{c}]$. The draws are independent of any of the other stochastic processes including the process leading to the initial queue order. Before being served, customers interact in an auction mechanism that allows them to trade places. Interacting in the auctions is assumed not to cost any (additional) time for the customers. Customer i 's utility from interacting in the auction is given by

$$U_i = \sum_{\substack{j=1 \\ j \neq i}}^N (P_{ji} - P_{ij}) - c_i w_i$$

where P_{lm} denote payments from customer l to customer m and w_i customer i 's total waiting time (i.e., time spent in the queue). We assume customers' service time to be equal to one time unit. Thus, if a customer is the k th to be served, she waits $k - 1$ time units in the queue, $k = 1, \dots, N$. We assume that all customers arrive before the server opens. A customer leaves the system after being served. We consider two auction mechanisms, the 'server-initiated auction' and the 'customer-initiated auction.' While our environment is

essentially static (all customers arrive before the server opens), we describe both auctions in such a way that they could be straightforwardly applied in a dynamic setting (where customers arrive while the server is busy).

Server-initiated auction. When idle, the server initiates an auction if two or more customers are in the queue. In this auction, each customer in the queue independently submits a bid. The server starts serving the customer who has submitted the highest bid. In the case of a tie, a fair lottery determines which customer gets served. This customer pays each of the r remaining customers a fraction $1/r$ of her bid. The winning bids are revealed to all customers. The losing bids are not revealed. Table 5.1 illustrates the rules of the server-initiated auction on the basis of a numerical example.

Table 5.1: Numerical example for the rules of the server-initiated auction

Initial queue order	First auction		Queue order	Second auction		Final queue order
	Bids	Transfers		Bids	Transfers	
A	18	+82	C			C
B	50	+82	A	15	+42	B
C	164	-164	B	42	-42	A

Notes. Three customers are in the queue when the server becomes idle. In the first auction, all three place a bid. In this case customer C submits the highest bid (164) and moves to position 1. His bid is distributed equally among the other two bidders. In the second auction, customer B places the highest bid (42) and moves to position 2. Customer B pays her bid to customer A.

Customer-initiated auction. Suppose there are $n \geq 1$ customers in the queue when a new customer arrives. The arriving customer is located at the end of the queue. She then trades places with the existing customers on the basis of the following algorithm:

1. $i \equiv n$.
2. Both the arriving customer and the customer directly in front of her independently submit a bid, which is denoted by b_{n+1}^i and b_i respectively.
3. If the customer in the queue in front of the arriving customer has submitted a bid $b_i > b_{n+1}^i$, the arriving customer remains in her current position and the process ends. Otherwise, go to step 4. (The bids are not revealed to any of the other customers.)
4. The arriving customer pays b_{n+1}^i to the customer in front of her. If $i = 1$, she stops. Otherwise, $i \leftarrow i - 1$. Return to step 2.

Table 5.2 contains a numerical example illustrating the rules of the customer-initiated auction.

Table 5.2: Numerical example for the rules of the customer-initiated auction

Initial queue order	First auction		Queue order	Second auction		Final queue order
	Bids	Transfers		Bids	Transfers	
A	76	+158	B			B
B	158	-158	A	100	0	A
			C	50	0	C

Notes. Three customers arrive in the queue. The first auction is initiated when customer B enters the queue. Customers A and B both submit a bid. As customer B places the higher bid, she swaps positions with customer A in return for a payment equal to customer B's bid. A second auction is initiated when customer C arrives. Both customer A, the second in line, and customer C submit a bid. Because the bid of the arriving customer is lower than the bid of the customer in front, there is no swap and, thus, there is no monetary transfer between the two customers.

As soon as the server completes serving one customer, it starts serving another, either the highest bidder (in the server-initiated auction) or the one at the front of the queue (in the customer-initiated auction). Note that both auctions are sequential games with incomplete information. We solve the games using the perfect Bayesian Nash equilibrium (henceforth: equilibrium). We obtain the following results. First of all, the server-initiated auction has a symmetric equilibrium. Let $B_n(c)$ denote the bid for a customer with waiting costs c in the case that n other bidders are in the queue.¹⁴

¹⁴ Proofs of propositions 1 and 3 and corollary 2 are relegated to the appendix.

PROPOSITION 1. Let $c_n^{(1)}$ represent the highest-order statistic among n independent draws from F , $n = 2, 3, \dots, N$. The following iteratively defined set of bidding functions constitutes an equilibrium of the server-initiated auction:

$$B_1(c) = \frac{1}{2} E \left\{ c_2^{(1)} \mid c_2^{(1)} \leq c \right\}$$

$$B_n(c) = \frac{n}{n+1} E \left\{ B_{n-1}(c_{n+1}^{(1)}) + c_{n+1}^{(1)} \mid c_{n+1}^{(1)} \leq c \right\}, n$$

$$= 2, 3, \dots, N-1.$$

In our experiment, we let the customers draw waiting costs from a uniform distribution. The following proposition establishes the resulting equilibrium.

COROLLARY 1. Suppose $F = U[0, \bar{c}]$ where $\bar{c} > 0$. Then

$$B_n(c) = \frac{nc}{3}, n = 1, 2, \dots, N-1$$

constitutes an equilibrium of the server-initiated auction.

Observe that in equilibrium all customers in the queue use the same strictly increasing bidding function for each position the server auctions. As a consequence, the highest bidder is always the customer with the highest waiting costs so that the bidders are served in order of waiting costs. The following result is then immediate.

COROLLARY 2. The server-initiated auction has an efficient equilibrium.

In contrast, for the customer-initiated auction, no efficient equilibrium exists. This result follows immediately from the

analysis by Gershkov and Schweinzer (2010) who show that in our setting no efficient individually rational and budget-balanced mechanism exists if individual rationality is with respect to the initial first-come, first-served order.

PROPOSITION 2. The customer-initiated auction does not have an efficient equilibrium.

Proof: Follows directly from Proposition 2 in Gershkov and Schweinzer (2010).

Proposition 3 illustrates this result by comparing equilibrium bids for the first position in the queue. It shows that a customer in position 1 at any point in the auction process bids more aggressively than the customer currently in position 2. As a consequence, for a non-zero mass of cost realizations, the arriving customer bids less than the first in line even if the arriving customer has higher marginal waiting costs. So, the two do not trade places, resulting in an inefficient queue order.

PROPOSITION 3. In any equilibrium of the customer-initiated auction, a customer in position 2 bids strictly less than the customer in position 1 conditional on the two having the same waiting costs $c > \underline{c}$.

The finding that the customer-initiated auction does not guarantee an efficient queue order is not surprising in light of the Myerson-Satterthwaite (1983) impossibility theorem. The theorem states that no efficient trade is feasible between a seller and a buyer if both are

incompletely informed about each other's value for the good owned by the seller and the range of possible buyer and seller valuations overlap. The impossibility result applies to the customer-initiated auction because the arriving customer is a potential buyer of the position in front of her and the range of values for two customers overlap.

5.3. Experimental design and hypotheses

5.3.1. Experimental design

We ran the computerized laboratory experiments at the Center for Experimental Economics and political Decision making (CREED) of the University of Amsterdam. Each session consisted of four parts. In all parts, participants interacted within groups of the same five participants (no re-matching). In the first part, participants interacted five times in either the server-initiated auction or the customer-initiated auction. In the second part, they interacted five times in the other auction mechanism. In part 3, the participants were asked to vote between the auction mechanisms played in the first two parts. Majority voting determined which of the two auction mechanisms was played in part 4, where we took the votes from all participants in a session together. In part 4, the participants interacted in the chosen auction mechanism and the final queue order determined when the participants could leave the laboratory. Thus, the last part incurred real waiting costs in contrast to induced waiting costs as in the first two parts.

Table 5.3: Experimental design and number of participants

		Arrival time	
		$t = 0$	$t \sim U[-4,0]$
		Simultaneous	Sequential
Order of auction mechanisms	Server –	20 (4)	20 (4)
	Customer		
	Customer –	30 (6)	25 (5)
	Server		

Note. Number of groups in parentheses.

We exploit a 2x2 between-subjects design where the treatments vary on two dimensions: the order of the auction mechanisms and the arrival process (see Table 5.3). In all treatments, the server initiates service at time 0 and service time is fixed at 1. Before entering the queue, participants drew their waiting costs per unit of time from the uniform distribution on the integer values from 0 to 100. All draws throughout the experiment were independent of each other and of any of the other stochastic processes. For the sake of comparison between the treatments, we kept the waiting cost draws constant across participant groups. We used the following two arrival processes. The first implements simultaneous arrivals: all customers arrive at time 0 and are put in a queue in random order. The second arrival process is a modification of Stein et al.’s (2007) sequential arrival protocol. All participants draw an arrival time according to the uniform distribution on the time interval $[-4,0]$. Upon arrival, each customer is located at the end of the queue and incurs waiting costs equal to the time she has to wait

until the server initiates service multiplied by her waiting costs per unit of time.

At the end of the experiment, we paid the participants their experimental earnings in the order determined in part 4. We left five minutes between paying each participant in the same participant group. As a consequence, the last student left the experiment 20 minutes after the first. By doing so, we induced actual waiting costs for the participants. Before we paid the participants, we asked them to fill out a questionnaire that included questions about background characteristics such as age, gender, and field of studies. In addition, the participants had to indicate on a seven-point Likert scale to what extent they considered the two auction mechanisms to be fair. Only when all participants in a session had finished the questionnaires did we start paying them.

At the start of the experiment, participants obtained a starting capital equal to 4,000 [3,000] ‘francs’ in the case of the sequential [simultaneous] arrival process. In all treatments, the exchange rate was 100 francs = €1. Earnings varied between €5.60 and €44.20, with an average of €20.97.¹⁵ We could conclude all sessions within two hours, including the 20 minutes the students at the end of the queue in part 4 had to wait.

5.3.2. *Hypotheses*

Our experimental design allows us to test several hypotheses. Our main theoretical finding is that the server-initiated auction has an

¹⁵ Participants earned on average €20.93 in the treatments with the simultaneous arrival process and €21.02 in the treatments with the sequential arrival process. There is no significant difference (Mann-Whitney, $U = 1031$, $p = 0.48$).

efficient equilibrium, in contrast to the customer-initiated auction. This result implies the following testable hypothesis.

HYPOTHESIS 1. *The server-initiated auction results in a more efficient outcome than the customer-initiated auction.*

Hypothesis 1 may be rejected if bidding behavior is ‘noisy’ in the sense that customers do not bid according to the same, strictly increasing bidding function. Consider the extreme case that the initial queue is already in the efficient order. Adding independent noise to the equilibrium bidding functions of the server-initiated auction implies that the actual service order may be inefficient. For the customer-initiated auction, the effect of adding independent noise may be more innocent than for the server-initiated auction when an arriving customer bids less aggressively than customers in front of her so that inefficient trade may be less likely to occur. As a consequence, noisy bidding behavior may imply that for relatively efficient initial queue orders, the customer-initiated auction is at least as efficient as the server-initiated auction so that hypothesis 1 is rejected.

In addition, as discussed in the introduction, the endowment effect and the sunk-cost effect may play a role in auctions that reallocate queuing positions. If an endowment effect is present, the alternative hypothesis is that a customer’s bid depends on her initial position in the queue. In the case of a sunk-cost effect, bids may depend on the arrival process because customers sink more costs before they get served in the case of a sequential arrival process than under a simultaneous arrival process. In contrast, the theory is based on the assumption that bidding behavior does not depend on either the

customers' initial positions or the costs customers sink before the server opens, which leads to the following hypotheses.

HYPOTHESIS 2. A customer's bids in the server-initiated auction do not depend on her initial position in the queue.

HYPOTHESIS 3. A customer's bids in the server-initiated auction do not depend on the waiting costs she sinks before the server opens.

The endowment effect might emerge differently in the customer-initiated auction. A salient difference between the server-initiated auction and customer-initiated auction is the ability to exercise position rights. Although customers in the server-initiated auction are assigned a rivalrous queue position, it is not possible to maintain the position because customers have no ability to defend it. On the contrary, customers in the customer-initiated auctions are assigned positions that can be defended. Customers can defend their position by submitting a very high bid if their position is contested without having to pay this bid. Thus, positions in the customer-initiated auction in terms of possession are 'stickier' than in the server-initiated auction. Reb and Connolly (2007) show that the endowment effect is more pronounced in cases of actual possession than in cases of perceived ownership. Therefore, it is more likely to find an endowment effect in the customer-initiated auction than in the server-initiated auction even though standard economic theory predicts that neither will appear.

The presence of position rights might also affect how sunk-costs affect bidding behavior. Baliga and Ely (2011) note that the sunk-cost effect can result in a willingness-to-pay that is higher or lower

than standard theory predicts. They argue that the sunk-cost effect is rational if decision-makers are assumed to have limited memory. The effect is either negative (the pro-rata effect) or positive (the Concorde effect) depending on how decision-makers use sunk costs as a signal about the future value of the ‘project.’ In the customer-initiated auction customers might perceive time spent in the queue as the price of obtaining position rights, while in the server-initiated auction customers cannot obtain these rights. Following Baliga and Ely (2011), time spent in the queue might lead to the Concorde effect in the customer-initiated auction and in the server-initiated auction it might lead to pro-rata effect. Both effects would be deviations from standard economic theory.

HYPOTHESIS 4. A customer’s bids in the customer-initiated auction, conditional on her current position and the history of play, do not depend on her initial position.

HYPOTHESIS 5. A customer’s bids in the customer-initiated auction, conditional on her current position and the history of play, do not depend on the waiting costs she sinks before the server opens.

Our final hypothesis concerns customers’ choice between the two auctions. In part 3 of the experiment we asked the participants to vote for one of the two auctions before they knew their actual position in the queue. Because the theory predicts that the server-initiated auction outperforms the customer-initiated auction in terms of efficiency gain, and because the efficiency gains are shared among the customers, we expect participants to prefer the server-initiated auction.

HYPOTHESIS 6. *The participants will vote for the server-initiated auction rather than the customer-initiated auction.*

5.4. Results

In this section, we present our experimental observations.¹⁶ Table 5.4 provides an overview of the main descriptive statistics. Because both auctions are budget balanced, comparing the participants' payoffs provides a first impression of the auctions' relative efficiency gains. Payoffs are higher in the server-initiated auction than in the customer-initiated auction for both arrival processes, suggesting that efficiency gains are higher in the former, in line with hypothesis 1. In the next section, we will study efficiency gains using a more refined definition. Moreover, in both auctions, bids are significantly higher if customers arrive before the start of the auctions (for both auctions, $p < 0.01$, two-tailed Mann-Whitney). This suggests that hypotheses 3 and 5 should be rejected, which could point to the presence of a sunk-cost bias. We look deeper into this in sections 5.4.2 and 5.4.3.

The descriptive statistics indicate that participants may have economic reasons to prefer the server-initiated auction over the customer-initiated auction. First, their average payoff tends to be higher in the server-initiated auction, albeit not significantly so (simultaneous arrival protocol: $p = 0.22$, sequential arrival protocol:

¹⁶ We find that bids in the first part are on average higher than in the second part. However, this effect is not dependent on the order of the auction mechanisms. Therefore, in our analysis we pool all data in parts 1 and 2. Our results are not qualitatively affected if the order of the auction mechanisms is controlled for.

$p = 0.78$, two-tailed Mann-Whitney). Moreover, the number of auctions per round is significantly higher in the customer-initiated auction than in the server-initiated auction ($p < 0.01$, two-tailed Wilcoxon Signed Ranks), which suggests that it will take less time to complete the auction process in the case of the server-initiated auction. As a result, the server-initiated auction might be preferred over the customer-initiated auction. However, we do not observe significant differences in terms of the number of auctions played per participant ($p = 0.11$, two-tailed Wilcoxon Signed Ranks). Moreover, when all customers arrive before the start of the auctions, both the difference between the highest and lowest payoff and the payoff variance are higher in the server-initiated auction (for both comparisons, $p = 0.05$, two-tailed Mann-Whitney), which suggests that customers might consider the customer-initiated auction more fair. However, if all arrive at the start of the auction, the opposite results obtain, although the differences are not statistically significant ($p = 0.71$ for the payoff range and $p = 0.60$ for the payoff variance, two-tailed Mann-Whitney). In section 5.4.4, we report participants' votes for the auction mechanisms as well as how fair they rate the two auction mechanisms.

Table 5.4: *Descriptive statistics*

	Server-initiated auction		Customer-initiation auction	
	<i>Arrival time</i>		<i>Arrival time</i>	
	$t = 0$	$t \sim U[-4,0]$	$t = 0$	$t \sim U[-4,0]$
Bid	82.77 (64.63)	98.77 (78.31)	57.08 (45.75)	65.19 (56.55)
# auctions per participant	2.80 (1.67)	2.80 (1.67)	2.70 (1.03)	2.64 (1.01)
# auctions per round	4.00 (0.00)	4.00 (0.00)	6.76 (1.98)	6.60 (1.91)
Payoff	-87.64 (110.32)	-188.96 (180.65)	-93.78 (126.89)	-190.64 (177.18)
Payoff Standard Deviation	115.16 (39.72)	185.89 (68.56)	126.99 (59.26)	167.65 (95.59)
Payoff Minimum	-211.14 (59.01)	-394.67 (90.95)	-254.54 (77.74)	-380.27 (132.41)
Payoff Maximum	77.96 (65.22)	62.60 (100.18)	58.40 (85.43)	35.87 (151.04)
Payoff Range	289.10 (107.10)	457.27 (172.71)	312.94 (145.56)	416.13 (204.08)

Notes. *Bid*: The average bid per participant per round in francs ($n = 950$). *# auctions per participant*: The average number of auctions per participant per round ($n = 950$). *# auctions per round*: The average

number of auctions per round ($n = 190$). *Payoff*: The average payoff per participant per round in francs ($n = 950$). *Payoff Variance*: The average standard deviation of the average payoff in a group per round ($n = 190$). *Payoff Minimum*: The average minimum payoff in a group per round ($n = 190$). *Payoff Maximum*: The average maximum payoff in a group per round ($n = 190$). *Payoff Range*: The average range between the minimum and maximum payoff in a group per round ($n = 190$). Standard deviations are in parentheses.

5.4.1. Efficiency gains

Firstly, we provide an overview of the two auctions' ability to improve the queue's efficiency. Customers enhance the queue's efficiency if a customer trades places with a customer behind her who has higher waiting costs. So, a natural measure of the queue's efficiency gain is the decrease in the sum of the customers' waiting costs after customers have traded places. More precisely, we define an auction's realized efficiency gain ΔE as

$$\Delta E \equiv \frac{W_{start} - W_{end}}{W_{max} - W_{min}}$$

where W_{start} [W_{end}] represents the sum of the customers' waiting costs when served according to the initial [final] queue order. For the sake of comparison between instances, we normalize an auction's efficiency gain by defining it as a fraction of the range of feasible efficiency levels, $W_{max} - W_{min}$, where W_{max} [W_{min}] stands for the highest [lowest] possible total waiting costs, i.e., the sum of the customers' waiting costs in the case that customers are

served in increasing [decreasing] order of waiting costs.¹⁷ Note that an auction's efficiency gain can be negative if the realized waiting costs are higher than the waiting costs that would have emerged if the customers had not traded places.

Table 5.5 shows that both auction mechanisms enhance queue efficiency on average. There were significantly more queues with a positive efficiency gain than a zero or negative efficiency gain in both auction mechanisms (server-initiated auction: Binomial, 62% positive, $p = 0.02$; customer-initiated auction: Binomial, 66% positive, $p < 0.01$). A single sample Wilcoxon Signed Rank test shows that the average realized efficiency gain is significantly greater than zero for both auctions (server-initiated auction: $Z = -5.57$, $p < 0.01$; customer-initiated auction: $Z = -6.06$, $p < 0.01$). Also at the group level, the efficiency gain is significantly greater than zero (server-initiated auction: $Z = -3.82$, $p < 0.01$; customer-initiated auction: $Z = -3.82$, $p < 0.01$).

We only find weak support for hypothesis 1. Table 5.5 shows that on average, the efficiency gain in server-initiated auction is equal to 0.33 while the average efficiency gain in the customer-initiated auction equals 0.28. So, queues using server-initiated auctions experience higher efficiency gains than queues using customer-initiated auctions. However, the difference is not statistically significant ($p = 0.24$, one-sided Mann-Whitney U test).

¹⁷ See, e.g., Goeree and Offerman (2002) for a similar measure of realized efficiency in auctions.

Table 5.5: Average efficiency gains

Auction mechanism	All	Low initial efficiency (< 0.50)	High initial efficiency (≥ 0.50)
Server-initiated	0.33 (0.05)	0.68 (0.03)	-0.18 (.03)
Customer-initiated	0.28 (0.04)	0.50 (0.03)	-0.06 (.02)
Difference	0.05 (0.06)	0.17 (0.04)	-0.12 (.03)
N (Queues)	190	114	76

Notes. Numbers represent the average efficiency gain (standard errors are in parentheses).

***/**/* Significant at the 1%/5%/10% level (two-sided Mann-Whitney U test)

Further analysis shows that the initial queue efficiency determines to what extent the mechanisms are able to enhance efficiency. In the last two columns of Table 5.5, we distinguish between queues with low and high initial efficiency, where initial efficiency is defined as

$$E_{start} \equiv \frac{W_{start} - W_{min}}{W_{max} - W_{min}}.$$

We find that server-initiated auctions are significantly more effective in increasing queue efficiency than customer-initiated auctions if the initial efficiency is low. In contrast, if the initial efficiency is high, queues using server-initiated auctions result on average in a significantly lower efficiency than queues using customer-initiated auctions. The regressions in Table 5.6 confirm that efficiency gains depend on the type of auction and initial efficiency. Customer-initiated auctions seem to be more rigid than server-initiated auctions, which is advantageous if the initial efficiency is high but impedes efficiency if this is low.

Our finding has the following intuitive explanation. Queues with a low initial efficiency can potentially gain more in terms of efficiency than queues with a high initial efficiency. Also, queues with a high initial efficiency risk decreasing in efficiency in the case of inefficient swaps. Both efficiency gains and efficiency losses are more likely to occur in the server-initiated auction than in the customer-initiated auction. The reason is that in contrast to the server-initiated auction, the customer-initiated auction protects position rights in the sense that the current position holder can retain her own position by submitting a high bid. In queues using the customer-initiated auction changes are expected to be less pronounced because incumbents are likely to block inefficient swaps.

Table 5.6: Estimation of efficiency gains

<i>Variable</i>	<i>Coefficient (S.E.)</i>	
Constant	0.59 (0.04)	***
Initial efficiency	-0.65 (0.05)	***
Auction mechanism (1 = Server-initiated, 0 = Customer-initiated)	0.25 (0.06)	***
Initial efficiency \times Auction mechanism	-0.38 (0.07)	***
<i>F</i>	220.13	***
<i>R</i> ²	0.78	
<i>N</i>	190	

Notes. OLS regressions with standard errors clustered at the group level.
***/**/* Significant at the 1%/5%/10% level

5.4.2. Individual bidding behavior in the server-initiated auction

In this and the next sections, we take a deeper look into individual bidding behavior in the two auction mechanisms to answer the question of why the two auction mechanisms do not differ significantly in terms of average efficiency gain. In this section we focus on the server-initiated auction. Standard economic theory predicts that the auction outcome is efficient because for each position customers bid according to the same bidding functions that are strictly increasing in waiting costs. Table 5.7 presents the results of five regressions on the bids submitted in the server-initiated auction. The estimated coefficients of the interaction term between waiting costs and the number of remaining other bidders are all significantly greater than zero and estimates range from 0.30 to 0.33, which is very close to the predicted value of 1/3 (see Corollary 1). However, the predicted intercept is zero while the

estimated intercepts are all significantly greater than zero, which implies systematic overbidding. More importantly, bidding is very noisy in the sense that the R^2 is only about 0.25. Indeed, participants are not even close to using the same bidding function, which explains why the auctions do not always render efficient queues.

To what extent could an endowment effect explain the noise observed in participants' bidding behavior? According to hypothesis 2, a customer's initial position in the queue should not affect bidding behavior. The hypothesis implies that bids do not correlate with a customer's initial position in the queue. However, Model II in Table 5.7 shows that the initial position significantly affects bidding behavior ($p = 0.03$). Specifically, bids tend be *lower* if the bid is placed on the initial position of the customer, which is quite the opposite of the endowment effect. Thus, this finding allows us to reject hypothesis 2, albeit not in favor of an endowment effect.

Table 5.7: Estimations of bids in the server-initiated auction

Model	Whole sample				Sequential arrival process			
	I	II	III	IV	V	VI	VII	VIII
Variable	Coefficient (S.E.)	Coefficient (S.E.)	Coefficient (S.E.)	Coefficient (S.E.)	Coefficient (S.E.)	Coefficient (S.E.)	Coefficient (S.E.)	Coefficient (S.E.)
Constant	24.69 (3.03) ***	25.72 (3.21) ***	24.73 (3.05) ***	26.44 (3.93) ***	34.94 (7.85) ***			
Waiting costs × # bidders left	0.30 (0.01) ***	0.30 (0.03) ***	0.26 (0.03) ***	0.33 (0.04) ***	0.33 (0.04) ***			
Bid on initial position		-5.35 (0.03) **						
Arrival process × # bidders left × Waiting costs			0.08 (0.05) *					
Arrival time							4.46 (1.85) **	
<i>F</i>	439.91 ***	221.16 ***	236.76 ***	238.59 ***	123.21 ***			
<i>R</i> ²	0.25	0.25	0.26	0.28	0.28			
<i>N</i>	1330	1330	1330	630	630			

Notes. OLS regressions with standard errors clustered at the group level. Arrival process is a dummy which equals 1 if and only if the observation concerns a sequential arrival process.
 ***/**/ Significant at the 1%/5%/10% level

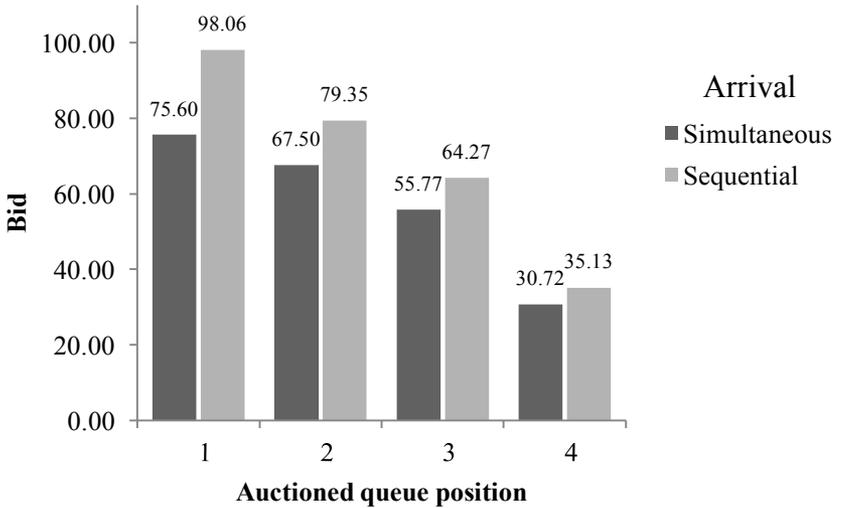
To identify potential sunk-cost effects, we test whether the arrival process affects bidding behavior. According to hypothesis 3, it should not because arrival costs are sunk at the time of bidding. Figure 5.1 suggests that we can reject this hypothesis as bidders submitted significantly higher bids on any position when arriving before the server opens than when all arrived at time 0. For example, in the case of a sequential arrival process, the average bid for the first position is significantly higher (+29.7%) than with a simultaneous arrival process ($p < 0.01$, two-tailed Mann-Whitney U test). Moreover, bids are significantly more likely to be higher than the equilibrium bid when customers arrived sequentially rather than simultaneously (sequential: 68.4%; simultaneous: 58.6%; Fisher's exact, $p < 0.01$). Figure 5.1 indicates that the observed sunk-cost effect is relevant for any of the auctioned positions. The regression analysis in Table 5.7, Model III, confirms that the arrival process has a significant effect on the steepness of the used bidding curves.

To have a further look into the sunk-cost effect, we restrict the sample to only bids in the treatment with a sequential arrival process in Table 5.7. Because arrival time is exogenously determined, it is possible to analyze how the magnitude of sunk costs affects bidding behavior. A sunk-cost effect emerges if participants' bids depend on their arrival time. The estimated coefficient of arrival time is significantly positive (see Table 5.7, Model V), i.e., the more waiting costs a customer sinks, the less aggressively she bids.¹⁸ This finding conforms to Baliga and Ely's (2011) pro-rate effect. However, explained variance as expressed

¹⁸ Recall that arrival time is a negative number in the interval $[-4,0]$. The absolute value of arrival time measures how long a customer has to wait before the server opens.

by the R^2 hardly increases compared to Model IV, which does not correct for arrival time.

Figure 5.1: Average bids in the server-initiated auction by position and arrival process



Note. Numbers represent average bid for each queue position in server-initiated auctions.

5.4.3. Individual bidding behavior in the customer-initiated auction

Similar to bidding behavior in the server-initiated auction, in the customer-initiated auction, the bids are significantly correlated with marginal waiting costs (see Table 5.8). We distinguish between attacking customers and defending customers. An attacking customer is a bidder who initially arrived at the back of the queue and bids to be able to swap with the customer in front of her, whom

we refer to as the defending customer. Inefficiencies emerge in the customer-initiated auction for two reasons. First, we find that attacking customers and defending customers adhere to different bidding strategies. Specifically, defending customers tend to bid consistently higher than attacking customers despite bidding according to a less steep bidding function. Second, only defending customers seem susceptible to a sunk-cost bias, which results in significantly higher bids if customers arrived relatively early. No such effect was found for attacking customers. These findings suggest that the sunk-cost effect is more salient in cases when customers can protect the ownership of their position.

An endowment effect in the customer-initiated auction would result in a higher willingness-to-accept among customers who are defending their initial position than standard economic theory would predict. Consequently, it is expected that the likelihood of an efficient outcome is lower if the auction involves a customer defending her initial position. An outcome is considered efficient if the attacker swaps positions with the defender if and only if the attacker's waiting costs are higher than the defender's. A Logit regression reveals that the likelihood of an efficient outcome is not significantly lower if the defender is defending her initial position ($B = -0.06$, $p = 0.76$). Therefore, in line with Hypothesis 4, we do not find evidence for an endowment effect in the customer-initiated auction.

Table 58: Regressions explaining bids in the customer-initiated auction

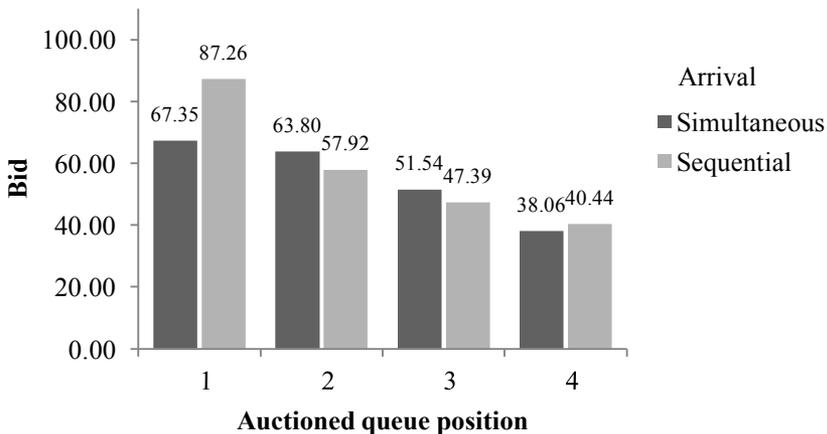
Model	Defending				Attacking			
	Whole sample		Sequential		Whole sample		Sequential	
	I	II	III	IV	III	IV	IV	
Variable	Coefficient (S.E.)							
Constant	34.01 (7.64) ***	21.27 (10.59) **	14.87 (4.53) ***	18.03 (7.04) **	14.87 (4.53) ***	18.03 (7.04) **	18.03 (7.04) **	
Waiting costs	0.53 (0.12) ***	0.59 (0.23) **	0.70 (0.07) ***	0.65 (0.13) ***	0.70 (0.07) ***	0.65 (0.13) ***	0.65 (0.13) ***	
Arrival time		-5.74 (2.91) **		-3.02 (3.32)		-3.02 (3.32)	-3.02 (3.32)	
F	32.80 ***	7.23 ***	140.02 ***	22.36 ***	140.02 ***	22.36 ***	22.36 ***	
R ²	0.05	0.05	0.18	0.13	0.18	0.13	0.13	
N	635	297	635	297	635	297	297	

Notes: OLS regressions with standard errors clustered at the group level. The dependent variable is the bid.

*** / ** / * Significant at the 1%/5%/10% level

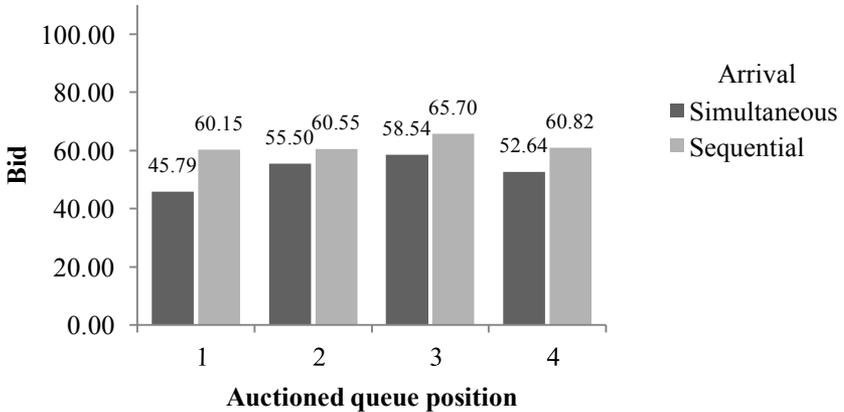
To study potential sunk-cost effects, we now focus on whether the type of arrival process affects bidding behavior in the customer-initiated auction. Figures 5.2 and 5.3 suggest that customers bid more aggressively in the case of a sequential arrival process. Like in the server-initiated auctions, average bids are higher for any of the positions auctioned under a sequential arrival process than under a simultaneous arrival process ($Z = -2.47, p = 0.01$, Mann-Whitney). In support of the sunk-cost bias and in conflict with hypothesis 5, we find that arrival time significantly affects bidding behavior in the case of a sequential arrival process (Table 5.8, Model II). Bids are increasing in the amount of waiting time before the server opens, which can be interpreted as evidence of the Concorde effect.

Figure 5.2: Defenders' bids in the customer-initiated auction by arrival process



Notes. $N = 635$. Numbers represent average bid for each queue position in customer-initiated auctions.

Figure 5.3: *Attackers' bids in the customer-initiated auction by arrival process*



Notes. $N = 635$. Numbers represent average bid for each queue position in customer-initiated auctions.

5.4.4. Preferred auction mechanism and actual waiting

In the third part of a session, participants were asked to vote for which auction mechanism to play in an experimental protocol that involved actual waiting. 63 percent of all participants voted for implementing the server-initiated auction in the last part of the experiment, which differs statistically significantly from 50 percent ($p < 0.01$, one-tailed binomial test). In fact, in all sessions a majority voted for the server-initiated auction. This finding provides some support to hypothesis 6, which states that customers will vote for the server-initiated auction rather than the customer-initiated auction.

Participants were asked at the end of the experiment to indicate for each auction mechanism to what extent trading positions is ‘fair’ on a five-point Likert scale. The average score for the server-initiated auction [customer-initiated auction] is 3.36 (SD = 1.07) [2.89 (SD = 1.11)]. The difference is statistically significant (Wilcoxon Signed Rank test, $Z = -3.30$, $p < 0.01$). Furthermore, participants who considered the server-initiated auction to be strictly more fair than the customer-initiated auction were also more likely to vote for the server-initiated auction (two-tailed Fisher’s exact test, $p = 0.02$). These findings suggest that fairness considerations partially underpin the preference for the server-initiated auction.

Table 5.9: Estimations of likelihood voting for server-initiated auction

<i>Model</i>	<i>I</i>	<i>II</i>	<i>III</i>
<i>Variable</i>	<i>Estimate</i> <i>(S.E.)</i>	<i>Estimate</i> <i>(S.E.)</i>	<i>Estimate</i> <i>(S.E.)</i>
Constant	0.30 (0.38)	-0.13 (0.42)	-0.09 (0.47)
Gender (1 = Female)	0.27 (0.45)	0.23 (0.46)	0.27 (0.47)
Payoff difference	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
Auction order (1 = Server-initiated first)	0.13 (0.45)	0.25 (0.46)	0.17 (0.47)
Arrival protocol (1 = Sequential)	0.14 (0.43)	-0.01 (0.45)	-0.07 (0.48)
Fairness		1.09 * (0.47) *	1.16 ** (0.49)
Perceived duration			0.11

			(0.48)
χ^2	1.94	7.58	8.14
Nagelkerke's R^2	0.03	0.11	0.11
Percentage correct	63.2%	68.1%	66.7%
N	95	94	94

Notes. Logit regression with standard errors between parentheses. Payoff difference stands for customer's payoff in the server-initiated auction minus her payoff in the customer-initiated auction. Fairness – a dummy variable where 1 stands for perceived fairness of the server-initiated auction – is strictly higher than that of the customer-initiated auction. Perceived duration – a dummy variable where 1 stands for perceived duration that is strictly higher than that of the customer-initiated auction. Similar results are found if perceived fairness and duration are weakly higher (Fairness: $B = 1.56, p = 0.01$; Perceived duration: $B = -0.79, p = 0.42$). In Model II and III one participant is excluded due to nonresponse. ***/**/* Significant at the 1%/5%/10% level

In Table 5.9 the likelihood of voting for the server-initiated auction is estimated. We find that only fairness considerations significantly affect the likelihood of voting for the server-initiated auctions. A participant's vote does not depend in a statistically significant way on differences in earned payoffs in the two auction mechanisms, the order of the auctions or the arrival protocol. The arrival protocol does not have a significant effect on the voting decision either, even though it has an effect on the payoff distribution between the two auction mechanisms as shown in Table 5.4. Participants might also vote for the server-initiated auction because fewer of such auctions are required per round ($M = 4.00, SD = 0.00$) compared to the customer-initiated auction ($M = 6.68, SD = 1.94$). However, perceived duration of a round is not associated

with the likelihood of voting for the server-initiated auction (Table 5.9, Model III).

We also analyze bidding behavior in this last part. According to Corollary 1, the optimal equilibrium bid in the server-initiated auction correlates positively with the number of other remaining customers. We find indeed that, on average, bids decrease by about 9.21 as the number of other bidders decreases ($p < 0.01$). The implied average waiting costs are about €0.28 per five minutes, or €3.32 per hour, which is about 85% of the gross minimum hourly wage for 18-year-old employees at the time of the experiment.

5.5. Conclusion

In this chapter, we have experimentally studied two auction mechanisms that allow customers to trade places in queues. In the server-initiated auction, the server sequentially auctions the right to be served next and pays all customers who remain in the queue an equal share of the winning customer's bid. In the customer-initiated auction, arriving customers iteratively offer money to customers in the queue in order to swap positions. We have used two novel experimental protocols to examine the behavioral properties of both auction mechanisms. One protocol implements induced waiting costs, which allows us to compare the two auction mechanisms in terms of efficiency gains. In the second protocol, participants could trade places in a queue where they had to wait before they could leave the lab. We applied this protocol to determine which auction mechanism participants would prefer in a context that involved actual waiting.

Our most important findings are the following. First of all, on average, the server-initiated auction and the customer-initiated auction perform equally well in terms of efficiency gain. Second, the participants indicated that they found the server-initiated auction a fairer mechanism than the customer-initiated auction. In a way, this result is surprising too, because the customer-initiated auction protects customers' initial positions in contrast to the server-initiated auction. Third, when voting between the two auctions, the participants tended to favor the server-initiated auction.

In both auctions we observe a sunk-cost effect but we find no evidence of an endowment effect. For the customer-initiated auction, the latter result may be surprising because incumbents are able to defend 'their' positions, which might have induced a sense of entitlements that is associated with the endowment effect.¹⁹ A possible explanation of why we are unable to find the endowment effect in either auction mechanism might be due to the fact that positions are intangible or that changing positions can only, at worst, delay access to the focal service but cannot lead to being deprived of access.

The contribution of this study to the extant literature is threefold. To our knowledge, this study is the first to experimentally study priority auctions in a queuing context. Although a large number of papers have studied such auctions theoretically, an empirical

¹⁹ Reb and Connolly (2007) find similar results in another context.

investigation was lacking.²⁰ As predicted in previous studies, we find that priority auctions can substantially increase a queue's efficiency. At the same time, we observe that substantial inefficiencies emerge even in theoretically efficient auction mechanisms.

The second contribution is that our study illustrates the importance of considering mechanisms that do not have an efficient equilibrium. It allows for the evaluation of mechanisms that might be considered in practice such as the customer-initiated auction. Interestingly, the customer-initiated auction mechanism improved efficiency on average as much as the server-initiated auction mechanism, while only the latter has an efficient equilibrium.

Third, our experimental design allows one to study endowment effects and sunk-cost effects in an environment involving queues. Endowment effects have been found in a large range of contexts (see, e.g., Knetsch 1989; Kahneman et al. 1990; Heyman et al. 2004; Manzur et al. 2016) including queues where customers discourage queue jumping (Mann 1969; Milgram et al. 1986; Oberholzer-Gee 2006; Helweg-Larsen and LoMonaco 2008). We add to this literature that endowment effects may be weak in environments where customers can trade positions. The sunk-cost effect has also been documented extensively in the empirical literature (e.g., Arkes and Blumer 1985; Phillips et al. 1991; Soman and Cheema 2001; Soman and Gourville 2001; Offerman and

²⁰ See, e.g., Kleinrock (1967), Lui (1985), Glazer and Hassin (1986), Rosenblum (1992), Mitra (2001), Afèche and Mendelson (2004), Kittsteiner and Moldovanu (2005), Gershkov and Schweinzer (2010), Kayı and Ramaekers (2010), and Yang et al. (forthcoming).

Potters 2005; Baliga and Ely 2011). This study is the first to study the sunk-cost effect by manipulating waiting costs for queued customers. Using an experimental protocol that determines time spent waiting in the queue before service starts, we have observed that such sunk costs induce customers to bid more aggressively on average compared to a setting where customers do not sink costs before the server opens.

We envision the following avenues for further research. First of all, how do the auctions perform when entry into the queue is endogenous? The increased efficiency of the queue may attract additional customers to the queue. On the one hand, this may be efficiency enhancing as more customers use the valuable service provided by the server. On the other hand, additional entrants impose a negative externality on other customers in that they may have to wait longer. In particular, some may enter the queue only to collect payments by other customers without having a genuine interest in the offered service. In addition, and relatedly, it would be insightful to test the auction mechanisms in practice using field experiments. In a field setting the question of endogenous entry could be naturally answered.

Relatedly, endogenous entry decisions may also affect the initial queue order. In our experiment we imposed an exogenous queue order, which may be relevant for a large range of settings where people arrive at random in a queue (e.g., waiting lists for parking permits). We leave settings where the entry decision is endogenous

(e.g., queuing for concert tickets) for further research.²¹ At first sight our analysis suggests that the customer-initiated auction should be preferred over the server-initiated auction if and only if the resulting arrival process implies relatively efficient queues. However, arrival decisions may also depend on the auction mechanism used, which will complicate the analysis.

This study has several implications for the management of queuing systems in the field. Queues are ubiquitous. They are typically very costly to avoid in the case of consistent demand or supply shocks and they may be the only socially acceptable way to ration scarce demand. The results from our experiment demonstrate that customers can benefit from auction mechanisms that allow them to trade places. We have shown that priority auctions can improve a queue's efficiency considerably while the customers retain all the gains from trade. In practice, many settings may exist where such mechanisms could be implemented, ranging from the allocation of houses, spots in daycare centers, and access to sport facilities to the short-term trading of landing and take-off slots in airports, repair services after a natural disaster, and server allocation in Internet hosting centers. Also, in physical waiting lines such auction mechanisms could be implemented when customers make use of apps on their smartphones that allow them to trade positions using an online platform. Furthermore, we find that customers prefer the server-initiated auction to the customer-initiated auction. Our

²¹ For instance, Barzel (1974) shows that both efficient and inefficient queue orders may emerge in a setting with endogenous arrival times, depending on the correlation between the opportunity cost of time and the value of the "prizes" that are awarded on a first-come first-served basis. See also Holt and Sherman (1982).

results suggest that companies interested in offering their customers the opportunity to trade places while waiting to get served should use the server-initiated auction rather than the customer-initiated auction.

5.6. Appendix: Proofs of propositions

Proof of Proposition 1. Suppose, for the moment, that B_n is a strictly increasing function for all $n = 1, \dots, N - 1$. Because after each auction, the winner's bid is revealed, the remaining customers can infer the winner's waiting costs. Therefore, if $n + 1$ customers are left, it is common knowledge that their c_i 's are drawn from F conditional on c_i being in the interval $[\underline{c}, \bar{c}_n]$ where \bar{c}_n are the waiting costs from the winner of the previous auction with $n + 2$ bidders. Let

$$G(x) \equiv \left(\frac{F(x)}{F(\bar{c}_n)} \right)^n$$

denote the cumulative distribution function of the highest-order statistic among n draws from F conditional on the draw being in the interval $[\underline{c}, \bar{c}_n]$. Define $k_n(c, x)$ as the expected costs of not being served in an auction with n remaining competitors for a bidder having waiting costs c , where x denotes the highest costs among her competitors.

A customer with cost parameter c pretending to have cost parameter \tilde{c} obtains expected utility

$$U_n(c, \tilde{c}, \bar{c}_n) = \int_{\tilde{c}}^{\bar{c}_n} \left(\frac{B_n(x)}{n} - k_n(c, x) \right) dG(x) - B_n(\tilde{c})G(\tilde{c})$$

where the first [second] term on the right-hand side refers to the case that the customer does not win [wins] the auction.

In equilibrium, for $c > \underline{c}$,

$$\begin{aligned}
& \left. \frac{\partial U_n(c, \tilde{c}, \bar{c}_n)}{\partial \tilde{c}} \right|_{\tilde{c}=c} \\
&= - \left(\frac{B_n(c)}{n} - k_n(c, c) \right) G'(c) - B_n(c)G'(c) \\
&\quad - B'_n(c)G(c) = 0 \Leftrightarrow \\
& B'_n(c)G(c) + \left(1 + \frac{1}{n} \right) B_n(c)G'(c) = k_n(c, c)G'(c) \Leftrightarrow \\
& B'_n(c)G(c)^{1+\frac{1}{n}} + \left(1 + \frac{1}{n} \right) B_n(c)G'(c)G(c)^{\frac{1}{n}} = k_n(c, c)G'(c)G(c)^{\frac{1}{n}} \\
& \Leftrightarrow \\
& B_n(c)G(c)^{1+\frac{1}{n}} = \int_{\underline{c}}^c k_n(x, x)G(x)^{\frac{1}{n}}dG(x) \Leftrightarrow \\
& B_n(c) = \frac{n}{n+1} \int_{\underline{c}}^c k_n(x, x)d \left(\frac{G(x)}{G(c)} \right)^{1+\frac{1}{n}} \\
& \quad = \frac{n}{n+1} \int_{\underline{c}}^c k_n(x, x)d \left(\frac{F(x)}{F(c)} \right)^{n+1}.
\end{aligned}$$

(Check the case $\underline{c} > 0$.) Note that $k_1(c, x) = c$ and $k_n(c, x) = -U_{n-1}(c, c, x) + c$ for $n = 2, 3, \dots, N-1$. Now, the proposition follows because $k_n(x, x) = -U_{n-1}(x, x, x) + x = B_{n-1}(x) + x$ for $n = 2, 3, \dots, N-1$. (It is readily verified that B_n is a strictly increasing function for all $n = 1, \dots, N-1$, which is the assumption we started with.)

Proof of corollary 2. The proof is by induction. Note that

$$B_1(c) = \frac{1}{2} E \left\{ c_2^{(1)} \mid c_2^{(1)} \leq c \right\} = \frac{c}{3}.$$

Now, fix $n = 2, 3, \dots, N - 1$ and assume that $B_{n-1}(c) = \frac{n-1}{3}c$. It is well-known that for $F = U[0, \bar{c}]$, $E \left\{ c_n^{(1)} \mid c_n^{(1)} \leq c \right\} = \frac{n}{n+1}c$. Therefore,

$$\begin{aligned} B_n(c) &= \frac{n}{n+1} E \left\{ B_{n-1}(c_{n+1}^{(1)}) + c_{n+1}^{(1)} \mid c_{n+1}^{(1)} \leq c \right\} \\ &= \frac{n}{n+1} E \left\{ \frac{n+2}{3} c_{n+1}^{(1)} \mid c_{n+1}^{(1)} \leq c \right\} \\ &= \frac{n}{n+1} \frac{n+2}{3} \frac{n+1}{n+2} c = \frac{n}{3} c. \end{aligned}$$

Proof of Proposition 3. Let $B_k(c_k)$ denote a customer's equilibrium bid as a function of her waiting costs c_k , where $k = 2$ refers to an arriving customer reaching position 2 and $k = 1$ to the current customer in position 1. According to a standard argument, both the arriving customer and the bidder in front of her use strictly increasing bidding functions in equilibrium. Without loss of generality, we may assume that at the boundaries, $B_1(\underline{c}) = B_2(\underline{c})$ and $B_1(\bar{c}) = B_2(\bar{c})$. Let $\Phi_k(b) \equiv B_k^{-1}(b)$ denote the inverse function of the bidding functions ($k = 1, 2$). Note that bidders need not only obtain utility from the auction itself, but also from later auctions when trading places with customers who arrive later. Let $U_1(c, x)$ denote a customer's expected additional utility she obtains when occupying the first position after the auction if her [the other customer's] waiting costs equal c [x]. $U_2(c)$ represents a customer's expected additional utility if her waiting costs are equal to c and she ends up in position 2 after the auction.

The arriving customer having waiting costs $c > \underline{c}$ solves

$$B_2(c) \in \operatorname{argmax}_b \int_{\Phi_1(b)}^{\bar{c}} U_2(c) dF_1(c_1) + \int_{\underline{c}}^{\Phi_1(b)} (c + U_1(c, c_1) - b) dF_1(c_1).$$

The first-order condition of the maximization problem is given by

$$F_1'(\Phi_1(B_2(c)))\Phi_1'(B_2(c))(-U_2(c) + c + U_1(c, \Phi_1(B_2(c))) - B_2(c)) - F_1(\Phi_1(B_2(c))) = 0,$$

which implies

$$B_2(c) = c - U_2(c) + U_1(c, \Phi_1(B_2(c))) - \frac{F_1(\Phi_1(B_2(c)))}{F_1'(\Phi_1(B_2(c)))\Phi_1'(B_2(c))}.$$

When defending her position, a customer having waiting costs $c > \underline{c}$ solves

$$B_1(c) \in \operatorname{argmax}_b \int_{\underline{c}}^{\Phi_2(b)} U_1(c, c_2) dF_2(c_2) + \int_{\Phi_2(b)}^{\bar{c}} (B_2(c_2) - c + U_2(c)) dF_2(c_2).$$

The first-order condition:

$$U_1\left(c, \Phi_2(B_1(c))\right) - B_1(c) + c - U_2(c) = 0$$

which implies that in equilibrium,

$$B_1(c) = U_1\left(c, \Phi_2(B_1(c))\right) + c - U_2(c).$$

Suppose that $B_1(c) \leq B_2(c)$. As both B_1 and B_2 are strictly increasing, $\Phi_2(B_1(c)) \leq \Phi_1(B_2(c))$ so that

$$\begin{aligned} B_1(c) &= U_1\left(c, \Phi_2(B_1(c))\right) + c - U_2(c) \\ &\geq U_1\left(c, \Phi_1(B_2(c))\right) + c - U_2(c) > B_2(c) \end{aligned}$$

which establishes a contradiction. Therefore, $B_1(c) > B_2(c)$ for all $c > \underline{c}$.

5.7. Appendix: Experimental instructions

Translated from Dutch:

General instructions

Welcome to this experiment! You can earn money in this experiment. The amount that you will earn depends on your decisions and the decisions of other participants in the same experiment. Your earnings are paid to you privately at the end of the experiment.

It is impossible for us to relate your name to your decisions. Therefore, your anonymity is guaranteed. Keep your decisions private. Talking with the other participants during the experiment is not allowed.

During the experiment you can gain and lose points. At the end of the experiment these points are exchanged for euros. 100 points is equal to €1.00.

At the beginning of the experiment you will receive a deposit of [starting capital] points. The points that you earn during the experiment are added to your deposit. The points you lose are subtracted from your deposit.

The experiment consists of four parts. The first and second parts consist of five rounds each. At the beginning of a round you and four others will be given a random position in a queue. You can change positions using auctions. The type of auction in the first part differs from the type of auction in the second part.

In the third part, you can vote for the type of auction that you prefer. The type of auction with the most votes will be used to determine the order of a queue. Your position in this queue will determine when you can leave the experiment.

Instructions for parts 1 and 2

[*Simultaneous arrival process*: This part consists of five rounds. In each round you will be placed with four others in a queue. Your starting position within the queue is determined randomly.]

[*Sequential arrival process*: This part consists of five rounds. In each round you will be placed with four others in a queue, where your position depends on your arrival time. Your arrival time equals the number of turns that you need to wait before the first customer is served. The arrival time is determined randomly. At the beginning of each round, you will find your arrival time and starting position on the screen.]

You can change positions using auctions. Your final position determines how many turns you will have to wait before being served. You will incur waiting costs for each turn that you have to wait. Waiting costs are subtracted from your deposit.

The customer in position 1 does not have to wait and, therefore, does not incur any waiting costs. The customer in position 2 has to wait one turn. The customer in position 3 has to wait two turns. The customer in position 4 has to wait three turns. And the customer in position 5 has to wait four turns.

Your waiting cost per turn is an integer between 0 and 100. This is also true for the other four participants in the queue. For each round the waiting costs are randomly drawn by the computer for all customers, where every value between 0 and 100 has the same likelihood to be drawn. In each round the waiting costs are independent from the waiting costs in the previous rounds and the waiting costs of other participants.

Example: Imagine that your waiting costs are equal to 10 and that at the end of the round your position is 5. You have to wait four turns. The total waiting costs for that round are: $10 \cdot 4 = 40$.

Server-initiated auction

In this part, you can bid on each position in the queue. The round starts with an auction for position 1. The winner is the customer with the highest bid. (In case of a tie, the computer will determine who wins using a fair lottery.) The winner gets position 1 and distributes his bid evenly among the other bidders. The next auction is for position 2. The winner is again the customer with the highest bid. The winner's bid is distributed evenly among the customers behind him or her. The customer in position 1 does not get anything. Positions 3 and 4 are auctioned the same way. Winners cannot participate in the remaining auction within the same round.

Example: Imagine that position 2 is auctioned. The customers with position 2, 3, 4 and 5 can place a bid. Imagine that the customer with position 3 places the highest bid: 75. This customer goes to position 2 while the customer in position 2 moves to position 3.

The three bidders, who now stand behind the customer in position 2, receive each $75/3 = 25$.

Test questions

Imagine that your current position is 5 and that you can bid on position 1. Your bid is 10 and among the four other bids, 20 is the highest bid. What is the outcome?

- You win the auction and pay 10.
- You lose the auction and receive nothing
- You lose the auction and receive 5 (correct answer)

Imagine that your current position is 5 and that you can bid on position 4. You bid 50 and the other remaining bidder bids 25. What is the outcome?

- You win the auction and pay 50 (correct answer)
- You lose the auction and receive 25
- You win the auction and pay nothing

Customer-initiated auction

In this part, every customer gets the chance to swap positions using an auction. The round starts with the customer in position 2 (W2) joining the customer in position 1 (W1). W1 and W2 both place a bid. If the bid of W2 is higher than the bid of W1, then W1 and W2 swap positions. If the bid is lower then there is no swap. (In case of a tie, there is a 50% chance of a swap.) If W1 and W2 swap positions, then W2 pays his or her own bid to W1.

Then the third customer (W3) joins the queue (on position 3). He or she places a bid on position 2. The customer currently in position 2 also places a bid. W3 wins if his or her bid is higher. In that case, W3 and the customer in position 2 swap places and W3 pays this customer his or her bid. If W3 moves to position 2 then an auction starts for position 1. If W3 loses the auction for position 2 then there is no auction for position 1. In a similar way, the fourth and fifth customers are able to move forward in the queue.

Example: A fifth customer (W5) joins the queue in position 5. The first possible swap is with the customer in position 4 (W4). Imagine that W5 bids 100 and W4 bids 50. Because W5 placed a higher bid, W5 and W4 swap positions. W4 receives 100 from W5. W5 is now in position 4 and W4 is in position 5. Subsequently, W5 has the opportunity to swap positions with W3, who is in position 3. Because W3 placed a higher bid, there is no swap. The round is now completed. The customers will be served in the current order.

Test questions

Imagine that your current position is 5 and that you bid on position 4. You bid 10 and the customer in position 4 (W4) bids 22. What is the outcome?

- You win the auction, pay 10 and swap positions
- You lose the auction, receive nothing, and positions are not swapped (correct answer)
- You win the auction, pay nothing, and positions are not swapped

Imagine that your current position is 1 and that you can bid on your own position. You bid 10 and the customer currently in position 2 bids 5. What is the outcome?

- You win the auction, pay 10, and swap positions
- You lose the auction, receive nothing, and positions are not swapped
- You win the auction, pay nothing, and positions are not swapped (correct answer)

Instructions for part 3

In this part you can vote for the type of auction that will be used in the next part to determine the queue for leaving the experiment. In the next part you are put in a queue with four others. These are the same participants with whom you interacted in parts 1 and 2. The starting positions in this queue are determined randomly. You can change positions using auctions. Your final position determines when you can leave the experiment. In this part, you do not pay for any waiting costs, but you will be required to wait longer depending on your position in the queue. Every turn takes 5 minutes.

The customer in position 1 does not have to wait and can leave the experiment right away. The customer in position 2 has to wait a single turn, which takes 5 minutes. The customer in position 3 has to wait two turns, which takes 10 minutes. The customer in position 4 has to wait three turns, which takes 15 minutes. And the customer in position 5 has to wait four turns, which takes 20 minutes.

The bids are still displayed in terms of points. 100 points is equal to €1.00.

As previously mentioned, in this part you do not pay any waiting costs. You do pay any winning bids from your deposit. Payments by other participants are added to your deposit. Your deposit is paid to you if it is your turn.

You can vote for the type of auction that was used in part 1 and for the type of auction that was used in part 2. The auction with the most votes of all participants in the laboratory will be used to determine the queue order. If both types get the same number of votes, then the auction type is picked randomly.