Submitted to
manuscript (Please, provide the manuscript number!)

On Customer (Dis)honesty in Unobservable Queues: The Role of Lying Aversion

Arturo Estrada Rodriguez, Rouba Ibrahim, Dongyuan Zhan
University College London, 1 Canada Square, London E14 5AB
arturo.rodriguez.18@ucl.ac.uk , rouba.ibrahim@ucl.ac.uk, d.zhan@ucl.ac.uk

Queues where people misreport their private information to access service faster are everywhere. Motivated by the prevalence of such behaviour in practice, we construct a queueing-game-theoretic model where customers make strategic claims to reduce their waiting time, and the Manager decides on the static scheduling policy based on those claims to minimize the expected delay cost in the system. We develop a lying aversion model where customers incur both delay and lying costs. We run controlled experiments to validate our modelling assumptions regarding customer misreporting behaviour. In particular, we find that people do incur lying costs, and that their misreporting behaviour does not depend on changes in waiting times, but rather on the scheduling parameters. Based on the validated lying aversion model, we study the equilibrium that arises in our game. We find that under certain conditions, the optimal policy is to use an *honour system* where service priority is given according to customer claims. We also find that it may be optimal to *incentivize more honesty* by means of an *upgrading policy* where some customers who claim to not deserve priority are upgraded to the priority queue. We find that the upgrading policy deviates from the celebrated $c\mu$ rule.

Key words: Scheduling policy, priority queues, strategic customers, lying aversion, behaviour in queues.

1. Introduction

Queuing systems where people make dishonest claims in order to access service faster are everywhere. For example, in telephone triage systems, it is known that patients routinely exaggerate symptoms to get a doctor's appointment sooner (Kirton et al. 2020) or an ambulance faster (Jones 2020). In the United Kingdom, there have been reports of people lying about their symptoms (BBC 2020a) or their employment status (BBC 2020b) in order to receive COVID-19 tests faster at home. In this case, it was possible for people to be dishonest with impunity because the online booking system was explicit on its reliance on public honesty rather than employment checks (Weaver and Proctor 2020). This is similar to the honour system used for the scheduling of COVID vaccinations in some countries; e.g., in the United States (NPR 2021a), where claims of eligibility are not verified since requiring additional checks would add barriers to access, and is therefore undesirable or logistically difficult to enforce (NPR 2021b). While some systems rely on unverifiable customer claims to schedule services (despite knowing that misreporting is prevalent), other systems opt to not rely on customer claims at all, due to such dishonesty. For example, a chat-bot symptom checker application was recently commissioned by the National Health Service (NHS) to aid in booking general practitioner (GP) appointments for patients. This initiative was dropped, after initial pilot trials, because patients admitted that they would exaggerate their symptoms (in the text exchange with the chat-bot symptom-checker) in order to see the GP faster (Heather 2017).

The above examples share similar characteristics. First, only customers know their private information, which gives rise to a clear incentive to misreport in order to begin service faster. As claims are not verified, it is difficult to punish or fine untruthful behaviour, which creates an opportunity to misreport with complete impunity. Second, queues in those settings are unobservable, i.e., customers make claims without seeing the other customers waiting, or directly observing other customers' claims. This is because the service requests themselves are made either online or through the phone. Third, those services typically involve a single (or rare) interaction with the queueing system, so there is a strong degree of anonymity for customers. Thus, reputation concerns do not play a role, and community enforcement of desirable behaviour, which may occur with physical queues, is not possible (Allon and Hanany 2012). Fourth, due to the nature of those settings, price discrimination, which is a standard operational control used to induce truthful incentive-compatible claims, is not possible. For example, healthcare services provided by the NHS are publicly funded and cannot be individually priced. Similarly, due to access concerns and the social nature of those systems, inspection and admission controls, which are also standard operational controls, may be undesirable. For example, during the COVID-19 pandemic in the United Kingdom, many supermarket chains implemented a groceries delivery service where priority slots were reserved for vulnerable people who were not able to leave their homes (Kaidan 2020). For this service, customer claims were actually verified against a list of vulnerable people provided by the government. However, recent reports have found that because of "poor data", it took too long to identify more than a third of the 2.2 million vulnerable people who were struggling to gain access to food (Syal 2021). Finally, while there is evidence of customer untruthfulness, clearly it remains that not everyone misreports. However, it is not clear what properties of the scheduling policy (i.e., the order in which to serve customers) drive customer misreporting behaviour, if at all, or help to mitigate it.

In settings like the ones described above, managers aim to make efficient scheduling decisions in order to minimize customer waiting, especially for those customers with the most urgent needs. However, it is not clear how customers will strategically respond to different scheduling decisions and, thus, how to use potentially false customer claims. Indeed, while the above examples share the same set of properties and presumably a similar objective, in practice different scheduling policies are observed, as illustrated above. This gives rise to our main research questions: (1) How should service systems make scheduling decisions in settings where customers may potentially make dishonest and unverifiable claims? (2) How does the misreporting behaviour of customers depend on the scheduling policy? (3) Can the Manager incentivize more honest claims by relying on a judicious scheduling policy?

To address those questions, we consider a priority queueing-game-theoretic model and run controlled experiments to validate our modelling assumptions regarding customer misreporting behaviour. In our game, the Manager's information about true customer types (either high or low) is based solely on customers' own claims. At the beginning of the game, the Manager defines and communicates a scheduling policy according to which customers will be served. We focus on scheduling policies where customers are routed probabilistically, based on their claims, to either the high or low-priority queues. The Manager seeks to minimize the total expected delay cost in the system. Based on the scheduling policy and the expected waiting times in the system, customers decide on their individual claims. Customers have the incentive to misreport if doing so leads to a shorter waiting time. Finally, based on the scheduling policy and customer claims, customers wait in a queue and eventually receive service. Waiting times in the system depend on customers' strategic claims and on the scheduling policy, which must be chosen optimally based on those claims. In turn, customer claims are, themselves, based on the scheduling policy. Due to this endogeneity, we must characterize the equilibrium in our game.

Contributions. To model misreporting behaviour in our game, we consider a lying aversion model in which customers incur lying costs when they misreport. To validate the existence of lying costs and to derive sensible modelling assumptions on customer misreporting behaviour, we run controlled experiments where people can misreport their private information in order to wait less in a virtual queue. In our experiment, we vary the waiting times and the scheduling parameters, and estimate the impact of those changes on the misreporting probability. To the best of our knowledge, we are the first to investigate misreporting behaviour when time is at stake, rather than money. Moreover, in our queueing setting, customers are routed probabilistically based on their claims, i.e., the outcomes of lying are random. Studying the effect that uncertainty has on lying aversion is a key contribution of our paper. Indeed, very few papers in the literature investigate how the lying behaviour changes in uncertain environments and none of these papers consider a queueing setting as we do here. Importantly, we provide experimental evidence that the majority of explainable lying behaviour can be attributed to the scheduling parameters, rather than the waiting times themselves.

Based on our experimentally-validated lying-aversion model, we theoretically study the equilibrium that arises in the system and the optimal scheduling policy. We find that, as long as there is some level of honesty in the system, customer claims are informative in equilibrium. Thus, it is always optimal for the Manager to rely, in some way, on customer claims in scheduling. This gives support to our motivating examples above where customer claims are actively sought despite the fact that misreporting is possible. In particular, we find that if customers' lying aversion is not sufficiently sensitive (in a sense to be made more precise later), then the optimal scheduling policy, which arises at equilibrium, is to give priority to customers based on their claims. However, if the lying aversion is sufficiently sensitive, then it is optimal to randomly upgrade some low claims to the high priority queue, because doing so incentivizes honesty in the system. Importantly, this upgrading policy deviates from the celebrated $c\mu$ rule, which prioritizes customers in decreasing order of their $c\mu$ index, where c denotes the delay cost and μ denotes the service rate. This is noteworthy because the $c\mu$ rule has been shown to be optimal with uncertain, yet exogenous, customer claims (Argon and Ziya 2009, Bren and Saghafian 2019).

The remainder of the paper is structured as follows. In §2, we review the relevant literature. In §3, we describe the model primitives and our queueing game. In §4, we describe the Manager's problem. In §5, we describe the customer's problem and present a lying-aversion model. In §6, we present an experimental study to validate several assumptions of the lying-aversion model. In §7, we derive the equilibrium in the game and describe the optimal scheduling policy. Finally, in §8, we draw conclusions.

2. Literature Review

This paper is related to four different streams of literature, which we describe below.

Queues with uncertain customer information. There is a rich queueing-theoretic literature which studies scheduling decisions in priority queues where true customer types are perfectly known to the Manager. In particular, the $c\mu$ rule has been shown to be optimal for delay cost minimization in various settings (Cox and Smith 1991, Van Mieghem 1995). In practice, the true customer types may not be perfectly known to the Manager. Van der Zee and Theil (1961) study how exogenous misclassification errors affect optimal scheduling decisions in the system. Argon and Ziya (2009) assume that the Manager receives a signal from each arriving customer, where the signal is the probability that the customer is of a certain type. Argon and Ziya (2009) show that the Highest Signal First (HSF) scheduling policy, which is consistent with the $c\mu$ rule, yields the lowest long-run average waiting cost among all finite-class priority policies. Bren and Saghafian (2019) consider the case in which the type of an arriving customer is known but the service rate for each type has to be dynamically learnt through a data-driven optimization approach. They also derive an optimal scheduling policy that is similar to the HSF policy. Singh et al. (2021) assume that customer signals are the output of a data-driven classifier, and propose an integrated approach where the classifier and the prioritization policy are jointly optimized. In line with this stream of literature, we also investigate scheduling decisions when customer true types are not known to the Manager. However, unlike those papers, we focus on settings where customers have private information about their own types, and where they strategically manipulate the signals (claims) that they send to the Manager.

Queues with strategic customers. There is a rich queueing-economics literature which focuses on settings where customer information is private and where market mechanisms such as pricing (Mendelson and Whang 1990, Afeche and Mendelson 2004), auctions (Kittsteiner and Moldovanu 2005), or bribing (Kleinrock 1967, Lui 1985) can be used to induce truthful reporting in the system; see Hassin and Haviv (2003) for a survey. Hu et al. (2021) investigate a setting where customers are strategic in deciding whether to disclose personal information to the service provider. Similar to this literature, we focus on modelling queues with strategic customers who are delay sensitive. However, contrary to those papers which resort to money transfers in order to differentiate customers, we study how lying aversion influences customers' strategic behaviour, and show how the Manager can differentiate between customers, to a certain extent, without the use of incentive-compatible pricing/auction mechanisms. Indeed, we focus on settings where money transfers are not applicable, and instead investigate how the scheduling policy helps to incentivize truthful reporting.

Several papers investigate how scheduling decisions affect strategic customer behaviour, as we do here (Afeche 2013, Afeche and Pavlin 2016, Yang et al. 2021, Yang 2021). In particular, some of these works find that giving partial priority, e.g., by means of reducing the gap between the expected waiting times of the different priority classes, may be useful to induce desirable customer behaviour. Our optimal scheduling policy is a partial priority policy as well. However, in our paper, this partial priority must be reached through upgrading low claims (and not by other means of manipulating the waiting times) because of customers' preferences for truth-telling.

Behavioural queues. Our work is broadly related to papers studying the behavioural foundations of queueing systems; see Allon and Kremer (2018) for background. Shunko et al. (2018) study the behavioural impact of queue design on worker productivity in service systems which involve human servers. Buell (2021) identifies the negative effects of last-place aversion in queues. Armony et al. (2021) develop a game-theoretic model to assess the performance of pooling when behavioural servers choose their capacities strategically. Kim et al. (2020) study admission decisions in queues using behavioural models and controlled experiments. Wang and Zhou (2018) study how the queue configuration affects human servers' service time in a field experiment. Ülkü et al. (2020) investigate the relationship between waiting time and subsequent purchase decisions. Luo et al. (2022) study how customers in observable queues form their completion costs based on their position in line, the number of people that have been served since they joined the line, and the service speed. Althenayyan et al. (2022) investigate how line-sitting and express lines affect customers' satisfaction and fairness perceptions.

While there is ample empirical evidence illustrating the important role of social norms and preferences in queueing systems, Allon and Hanany (2012) is, to the best of our knowledge, the only work that investigates queueing intrusions with a formal mathematical model. In particular, Allon and Hanany (2012) show that the common observation of “queue jumping” can be part of social norms and can be explained on rational individual grounds. We depart from Allon and Hanany (2012) in three fundamental ways. First, Allon and Hanany focus on service systems with observable queues where the Manager is not involved in the way in which the queue is managed. Thus customers, through community enforcement, regulate the queue. In contrast, we focus on unobservable queueing systems where the Manager is in charge of controlling the queue, and where customers do not see each other and are unable to prevent intrusions. Second, unlike Allon and Hanany (2012), we model customer aversion to being untruthful and study its operational implications. Third, our methodological approach is different, as we conduct controlled experiments to shed light on the untruthful behaviour of customers in queues.

Misreporting behaviour and lying costs. In recent years, a fast-growing literature across economics, psychology, and sociology has begun to study how people misreport their private information; see Rosenbaum et al. (2014) for a survey and Abeler et al. (2019) for a meta-analysis. Overall, the literature strongly shows that people exhibit lying aversion (Abeler et al. 2019) as not everyone misreports even under conditions of complete anonymity and impunity. In particular, the experimental paradigm in Fischbacher and Föllmi-Heusi (2013) is the most widely adopted in the literature to investigate misreporting behaviour. In this paradigm, participants privately observe the outcome of a random variable, report this outcome, and receive a monetary payoff proportional to their report. Similarly, we adopt this experimental paradigm to study how participants misreport in a queueing setting. However, in contrast to this literature which uses *money* to incentivize participants to misreport their private information, in our experimental investigation, people have the incentive to misreport in order to shorten their *waiting time* in queue. Also, the literature has almost exclusively focused on settings where outcomes are certain, and until recently the literature has been largely silent on how lying costs extend to uncertain environments. In contrast, in our experimental investigation, outcomes associated with lying or telling the truth may be uncertain due to the scheduling policy. Similar to our work where outcomes are uncertain, Celse et al. (2019), Dugar et al. (2019), Steinel et al. (2022) explore, under different risky environments, how lying changes when its consequences are not certain and they do not consider a queueing setting. Contrary to this, we investigate how lying changes when the consequences of truth-telling (due to the upgrading probability) are random. Finally, Özer et al. (2011) also models lying costs, however, their paper focuses on forecasting in supply chains, unlike our focus in this paper.

3. Model Primitives

We consider an $M/M/1$ queueing system where customers arrive according to a Poisson process with rate λ, and where service times are independent and identically distributed exponential random variables with parameter $1/\mu$. An arriving customer has type X, where X takes value H with probability p_H, and value L with probability $p_L = 1 - p_H$. Customers are delay sensitive, and a customer of type $X = x$ has a per-time-unit waiting cost c_x. We assume that $c_H > c_L$. The traffic intensity is $\rho = \lambda/\mu$, and we assume that the system is stable, i.e., that $\rho < 1$.

At the beginning of the game, the Manager commits to and communicates a claim-based scheduling policy according to which customers will be served. Customers have private information about their type X and, upon arrival, make a claim $Y = y \in \{H, L\}$ about their type. Because the Manager does not know the true customer types, and because the scheduling of customers depends on their claims, customers may be untruthful in their claims, i.e., we may have $X \neq Y$. While the Manager cannot observe individual customer types, he can correctly anticipate the customers' aggregate claiming behaviour for a given scheduling policy. Finally, based on the scheduling policy and on customer claims, customers wait in queue and, eventually, receive service. Waiting times in the system depend on customer claims and on the scheduling policy, which must be chosen optimally based on those claims. In turn, customer claims are, themselves, based on the scheduling policy which is announced by the Manager a priori. That is, the dependence between waiting times, customer claims, and the scheduling policy is endogenously determined.

4. The Manager's Problem

The Manager commits to a scheduling policy according to which customers will be served. The Manager's objective is to minimize the expected waiting cost in the system. We investigate static work-conserving two-class priority policies, where customers are assigned to a priority class with expected waiting time W_y (not including service time) based on their claims $y \in \{H, L\}$.

To find the optimal scheduling policy, it is customary to use the *operationally achievable* method (Coffman Jr and Mitrani 1980), which focuses on identifying the optimal waiting times, W_y, factoring in customer decisions, while abstracting away from specific scheduling policies that induce those waiting times. Once the optimal W_y are identified, one can consider different policy implementations and calibrate the relevant parameters consistently. Importantly, it is implicitly assumed, when using this method, that customer decisions do not depend on the actual scheduling policy implementation - a reasonable assumption with rational customers. However, when behavioural factors are considered, as we do here, different scheduling policy implementations, corresponding to the same expected waiting times, may lead to different customer behaviours. In particular, customer misreporting behaviour has been found to be insensitive to changes in the benefits associated with dishonesty, but highly conditional and susceptible to contextual factors (Rosenbaum et al. 2014, Abeler et al. 2014), e.g., factors that increase the salience of honesty standards, factors that allow for moral justifications (Mazar et al. 2008), or factors that allow for reactions and interpretations to objective risk (Dugar et al. 2019). This is particularly relevant in our investigation of optimal scheduling policies as different implementations may vary in their contextual factors. For example, policies that schedule the service probabilistically based on customers' decisions (Yang et al. 2021) allow for different customer interpretations of objective risk, while policies that schedule the service deterministically (Yang 2021) do not. In Section 6.3, we provide experimental evidence substantiating that customer misreporting behaviour is driven by the actual scheduling implementation parameters alone, rather than by changes in the waiting times W_y. Accordingly, to derive managerial prescriptions that take into account customer behaviour, we must focus on a specific class of scheduling policies.

4.1. Scheduling Policy Implementation

We consider the non-preemptive version of the α*-policy* (Yang et al. 2021), since preemptive policies are difficult to implement in practice. Based on customer claims, the Manager assigns customers to a queue $K = k \in \{1, 2\}$ with expected waiting times W_1 (priority queue) and W_2 (regular queue). These waiting times arise endogenously and are given by

$$W_1 = \frac{\lambda}{\mu(\mu - \lambda_1)} \quad \text{and} \quad W_2 = \frac{W_1}{1 - \lambda/\mu},$$

where λ_1 is the rate of arrivals of customers who are assigned to the priority queue at equilibrium. Customers are served in a first-come-first-served (FCFS) fashion within each queue, and customers in the priority queue are given non-preemptive priority over customers in the regular queue. In particular, based on customer claims, the α-policy assigns customers to the priority queue with probabilities $\alpha_y = \mathbb{P}(K=1|Y=y)$. That is, upon arrival, customers who claim H (L) are assigned to the priority queue with probability α_H (α_L), and to the regular queue with probability $1-\alpha_H$ ($1-\alpha_L$). Thus, the Manager's problem is to select the *routing probabilities* $\boldsymbol{\alpha} = (\alpha_H, \alpha_L)$ in order to minimize the total expected steady-state waiting cost in the system:

$$\underset{\boldsymbol{\alpha}}{Min} \sum_{k\in\{1,2\}} \sum_{x\in\{H,L\}} \lambda_{k,x} c_x W_k, \tag{1}$$

where $\lambda_{k,x}$ is the equilibrium rate of arrivals of customers of true type x who are assigned to priority queue k. The Manager cannot observe the true type x of each customer, but can determine $\lambda_{k,x}$ based on the known equilibrium customer behaviour. Specifically, the arrival rate $\lambda_{k,x}$ is obtained by conditioning on customer claims as follows for $x \in \{H,L\}$ and $k \in \{1,2\}$:

$$\lambda_{k,x} = \lambda \cdot \mathbb{P}(X=x) \cdot \sum_{y\in\{H,L\}} \mathbb{P}(K=k|X=x, Y=y)\mathbb{P}(Y=y|X=x),$$

where $\mathbb{P}(K=k|X=x,Y=y) = \mathbb{P}(K=k|Y=y)$ since the assignment of claims to priority queues is independent of $X=x$, conditional on $Y=y$. Also, $\mathbb{P}(Y=y|X=x)$ captures a typical x type customer's equilibrium claiming behaviour. In §7, we describe the customer claiming behaviour and the optimal scheduling policy that arise in equilibrium.

We note that the α-policy is general in the sense that different parametrizations lead to conceptually different policies. For example, $(\alpha_H = 1, \alpha_L = 0)$ corresponds to a policy where customers with H claims are given *full priority*[1]. In general, $(1 \geq \alpha_H > \alpha_L \geq 0)$ correspond to policies that give *partial priority* to customers with H claims. For example, $(1 > \alpha_H > \alpha_L = 0)$ corresponds to a policy where customers with H claims are given partial priority by means of *downgrading* some H claims to the regular queue. Also, $(1 = \alpha_H > \alpha_L > 0)$ corresponds to a policy where customers with H claims are given partial priority, albeit by means of *upgrading* some L claims to the priority queue. Finally, $(\alpha_H = 1, \alpha_L = 1)$ corresponds to the FCFS policy. Note that, as above, we can also let $\alpha_L > \alpha_H$ to give full or partial priority to customers with L claims.

[1] With the α-policy, expected waiting times based on claims are given by $W_y = \alpha_y W_1 + (1-\alpha_y)W_2$. Formally, customers who claim H are given full priority whenever $W_H = \frac{\lambda}{\mu(\mu-\lambda_H)}$ and $W_H < W_L$. Also, customers who claim H are given partial priority whenever $W_H > \frac{\lambda}{\mu(\mu-\lambda_H)}$ and $W_H < W_L$. The same logic applies for L claims.

5. Customer Problem: A Lying Aversion Model

It is well established in the literature that when dishonesty is not sanctioned or remains anonymous, the decision of whether or not to lie poses a conflict between self-interest and self-concept (Fischbacher and Föllmi-Heusi 2013, Mazar et al. 2008, Rosenbaum et al. 2014, Abeler et al. 2019). That is, people have an incentive to lie to obtain a material gain (e.g., reducing their waiting cost) but they, concurrently, incur a lying cost because misreporting is intrinsically costly. To capture this tension, we model the behavioural process which leads to individual customer claims as follows. A customer of true type $x \in \{H, L\}$ makes a claim $y \in \{H, L\}$ to minimize a total expected cost equal to the sum of the expected delay cost and the intrinsic lying cost:

$$\underset{y \in \{H,L\}}{Min} \; c_x W_y + \theta \ell(x, y, \boldsymbol{\alpha}, \mathbf{W}), \tag{2}$$

where $c_x W_y = c_x(\alpha_y W_1 + (1-\alpha_y)W_2)$ is the expected waiting cost for customers of type x who claim y and $\theta\ell(x, y, \boldsymbol{\alpha}, \mathbf{W})$ is the expected lying cost of a misreporting customer, for $\boldsymbol{\alpha} = (\alpha_H, \alpha_L)$ and $\mathbf{W} = (W_1, W_2)$. Because customers are delay sensitive, they have the incentive to misreport their types if doing so leads to a shorter expected waiting time. However, a growing body of research has consistently found that people present lying aversion (Gneezy et al. 2013): People often refrain from misreporting their private information, even when they can do so anonymously and with complete impunity (Rosenbaum et al. 2014). Such lying aversion coalesces from a myriad of idiosyncratic factors such as moral or religious reasons, self-image concerns, or unwillingness to deviate from socially-acceptable behaviour (Abeler et al. 2014, Gibson et al. 2013).

To model intrinsic lying costs, we assume that customers experience no lying cost when they claim their true types, i.e., $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) = 0$ for $y = x$, and incur a non-negative lying cost whenever they misreport, i.e., $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) \geq 0$ for $y \neq x$, where such lying cost is a function of (potentially) the routing probabilities and waiting times. In §5.1, we present several sensible specifications for the lying cost, based on the extant literature. Consistent with empirical research which suggests that individuals exhibit heterogeneous lying costs (Gibson et al. 2013, Rosenbaum et al. 2014, Abeler et al. 2014, 2019), we assume that customers are endowed with a lying aversion, Θ, which is random. Each customer draws a random variate, $\Theta = \theta$, which represents their individual lying aversion, where θ is independently drawn from a distribution Φ with density ϕ over some interval $[0, \bar{\theta}]$. While individual θ values are private information to customers, we assume that the Manager knows the distribution of Θ.

Finally, we emphasize that, as per our motivating examples in the introduction, we focus on unobservable queueing settings, where there are no reputation concerns and where the service involves an anonymous, one-shot, interaction with the Manager. Since the queue is unobservable, customers do not see others making false claims and jumping the queue. We envision that, with observable queues, directly observing the behaviour of others may shape a customer's lying cost, i.e., we may have reputation-based lying costs[2]. We defer the study of lying behaviour in observable queueing setting to future research.

5.1. Sensible Lying-Cost Specifications

In the extant literature, there are two main approaches to model the intrinsic lying cost[3]: (1) A fixed lying cost (DellaVigna et al. 2016, Khalmetski and Sliwka 2019), and (2) a lying cost which increases in the linear (or strictly convex) *material benefit* derived from misreporting (Duch et al. 2021, Gneezy et al. 2018, Kartik 2009). In our setting, a fixed lying cost would be $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) = d$ for $y \neq x$, where d is some constant. A lying cost which increases in the convex *material benefit* derived from misreporting would be $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) = (c_x(W_x - W_y)^+)^r$, for $r \geq 1$. This expression captures the material benefit as the difference between expected waiting costs given a claim. We

[2] For example Dufwenberg and Dufwenberg (2018) present a model where people derive disutility in proportion to the amount by which they are perceived to misreport, even if they are actually honest. That is, in observable queues, our assumption that $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) = 0$ for $y = x$ does not necessarily hold.

[3] The other prominent classes of models found in literature that do not focus on intrinsic lying costs are: social norms models, where lying costs depend on the perception of how others behave, e.g., Abeler et al. (2019); and reputation models, where lying costs depend on what others think about the honesty of the individual, e.g., Gneezy et al. (2018), Dufwenberg and Dufwenberg (2018), Khalmetski and Sliwka (2019). See Abeler et al. (2019) for a survey on variations within each class: for example, the social norms class includes models based on conformity and inequality aversion; and the reputation class includes models based on reputation for honesty and reputation for not being greedy.

note that, in the operations management literature, Özer et al. (2011) propose a linear size of the lie specification in terms of the extent to which information is misreported. Ultimately, our objective is to identify to what extent the predictions under those sensible lying-cost specifications, hold with human decision-makers.

5.2. Misreporting Behaviour

Individual customer types and lying costs are *not observable*, and only the aggregate misreporting probability can be determined by the Manager. This is consistent with our motivating examples where individual customer misreporting cannot be directly observed. In such settings, studying how the misreporting probability depends on a given set of *fixed* values for the waiting times and routing probabilities, i.e., deriving the *best-response* misreporting probability function, allows us to identify which lying-cost specification, from the ones introduced above, is consistent with experimental data. Thus, we will focus our experimental efforts here on investigating the best-response misreporting probability function. This circumvents many experimental complexities associated with studying queueing steady-state equilibria (Allon and Kremer 2018).

Consider a tagged x type customer with lying aversion parameter θ' who faces, upon arrival, claim-based expected waiting times W_y and routing probabilities α_y. Based on (2), the expected cost of that tagged customer is $c_x W_x$ for a honest claim $y = x$ and $c_x W_y + \theta' \ell(x, y, \boldsymbol{\alpha}, \mathbf{W})$ for a dishonest claim $y \neq x$. That tagged customer selects the claim y which minimizes her expected cost. First, notice that the tagged customer will never misreport if there is no incentive to do so (i.e., $W_x < W_y$ for $y \neq x$) since $c_x W_x < c_x W_y + \theta' \ell(x, y, \boldsymbol{\alpha}, \mathbf{W})$. In contrast, if there is an incentive to misreport (i.e., $W_x > W_y$ for $y \neq x$), the tagged customer misreports whenever her lying aversion parameter θ' is sufficiently small[4]:

$$\theta' \leq \frac{c_x(W_x - W_y)}{\ell(x, y, \boldsymbol{\alpha}, \mathbf{W})} = \frac{c_x(\alpha_y - \alpha_x)(W_2 - W_1)}{\ell(x, y, \boldsymbol{\alpha}, \mathbf{W})} \quad \text{for} \quad y \neq x.$$

[4] Recall that $W_y = \alpha_y W_1 + (1 - \alpha_y) W_2$ for $y \in \{H, L\}$. Therefore, the difference in expected waiting times given a claim is $W_x - W_y = (\alpha_y - \alpha_x)(W_2 - W_1)$ for $y \neq x$.

Since individual lying aversions, i.e., specific θ values, are not observable, it follows that a typical x type customer misreports with probability:

$$\mathbb{P}(Y=y|X=x)=\Phi\left(\frac{c_x(\alpha_y-\alpha_x)(W_2-W_1)}{\ell(x,y,\boldsymbol{\alpha},\mathbf{W})}\right) \quad \text{for} \quad y\neq x, \tag{3}$$

where we recall that Φ is the CDF of the lying aversion Θ. The best-response misreporting probability is given by (3). This expression motivates our experimental design. In §6, we run controlled experiments to investigate how changes in pre-specified values for the waiting times and routing probabilities affect the misreporting probability in (3). The misreporting probability is readily measurable in a controlled experimental environment where subjects face an individual decision task. Based on the aforementioned potential lying-cost specifications, i.e., fixed cost, linear cost in the material benefit, and strictly convex cost in the material benefit, the best-response misreporting probability is, respectively, given by:

$$\text{[fixed]} \quad \mathbb{P}(Y=y|X=x)=\Phi\left(\frac{c_x(\alpha_y-\alpha_x)(W_2-W_1)}{d}\right) \quad \text{for} \quad y\neq x, \tag{4}$$

$$\text{[linear]} \quad \mathbb{P}(Y=y|X=x)=\Phi\big(1\big) \quad \text{for} \quad y\neq x, \tag{5}$$

$$\text{[strictly convex]} \quad \mathbb{P}(Y=y|X=x)=\Phi\left(\frac{1}{(c_x(\alpha_y-\alpha_x)(W_2-W_1))^{r-1}}\right), \text{with } r>1, \quad \text{for} \quad y\neq x. \tag{6}$$

We see in (4) that, with a fixed lying cost, misreporting would increase in both the difference in routing probabilities $\Delta\alpha=\alpha_y-\alpha_x$ for $y\neq x$ and in the difference in waiting times between queues $\Delta W=W_2-W_1$. In contrast, in (5), we see that with a lying cost that increases linearly in the material benefit, changes in routing probabilities $\Delta\alpha$ and waiting times between queues ΔW do not affect misreporting. Finally, in (6), if the lying cost increases convexly in the material benefit for $r>1$, then the probability of misreporting decreases in both $\Delta\alpha$ and ΔW. In the following section, we experimentally investigate the dependence of the misreporting probability on $\Delta\alpha$ and ΔW. This allows us to assess which lying-cost model, if any, is consistent with the data.

6. Experimental Investigation

In this section, we experimentally study the direction of changes in human decisions along the two dimensions of the scheduling policy environment: The routing probabilities and the waiting times.

In particular, we study how changes in both $\Delta\alpha$ and ΔW affect the best-response misreporting probabilities associated with subject decisions. This study is based on the well-known experimental procedure proposed by Fischbacher and Föllmi-Heusi (2013) to investigate the extent to which people misreport their private information, hereafter referred to as the FFH paradigm. In the FFH paradigm, participants *privately* observe the outcome of a random variable, e.g., the roll of a die. They are then asked to report that outcome and, subsequently, receive a payoff depending on their reported claims. While dishonesty cannot be verified at the individual participant level (because the true outcome is unknown to the experimenter), one can make inferences about aggregate participant lying behaviour, because the probability distribution of the random variable in the experiment is known to the experimenter. Importantly, participants can misreport their privately observed outcomes with absolute impunity and anonymity, making this design appropriate for investigating intrinsic lying costs - it allows for the exclusion of other confounding factors, such as reputation costs, negative externality towards others, perception by others, or fear of getting caught. The FFH paradigm is the most widely adopted in the literature: it has been used in over 90 studies involving more than 44,000 subjects across 47 countries (Abeler et al. 2019). Finally, several studies have found that the observed behaviour in the FFH paradigm represents a good metric for honesty as it correlates strongly with real-life behaviour (Gächter and Schulz 2016, Dai et al. 2018).

6.1. Experimental Procedure

We conduct an online virtual queueing experiment that represents a one-shot individual decision-making situation where participants are told the waiting times in two queues and the routing probabilities to these two queues, depending on their claims. In other words, scheduling is in accordance with the α-policy. Participants are randomly assigned to one of nine different experimental conditions in a between-subject fashion (see Table 1) where we vary the differences in the routing probabilities, $\Delta\alpha$, and the waiting times between queues, ΔW. Consistent with our unobservable queueing setting and motivating examples in the introduction, participants cannot observe the dynamic state of the queue, i.e., the number of participants in line, or the real-time behaviour of other participants. They only have information about the waiting times in each queue and the routing probabilities.

Based on the FFH paradigm, participants are asked to privately roll a six-sided die and to record the outcome on a piece of paper. To avoid potential confounding factors, participants are instructed to roll a die at home or, alternatively, to use Google's virtual die[5]. Indeed, as described above, to study intrinsic lying costs it is important that the experimenters do not observe the die rolls. After participants roll the die and write down the outcome on a piece of paper, they are presented with the waiting times and routing probabilities according to their randomly assigned condition. Participants are instructed that those who claim the number 5 will wait in the Short queue with probability α_H, and those who claim any other number will wait in the Short queue with probability α_L. We assume that $\alpha_H > \alpha_L$ in all experimental conditions, so that participants have the incentive to claim the number 5 to reduce their expected waiting time. Accordingly, participants report their private die roll. Participants are also asked, as part of the instructions for the experiment, some questions related to the α-policy implementation. Participants are able to finish the instructions section only if those questions are correctly answered. Moreover, before they roll the die and report any number, they are placed in a practice queue for 2 minutes. This ensures that participants understand what it feels like to wait in the Short virtual queue (in all our experimental conditions, the waiting time of the Short queue is 2 min). After participants report a number, they are assigned to a queue based on the α-policy, wait in queue, and finally answer two simple questions relating to their experience in the queue, which concludes the experiment. Throughout the experiment, there is no mention of lying, honesty, or any related concepts.

Participants and exclusions. We recruited participants from the Amazon Mechanical Turk (MTurk) platform. Participants with at least 0.95 HIT approval ratio (proportion of completed tasks) were recruited to take part in our queueing experiment. Participants were instructed that

[5] At the following link `https://www.google.com/search?q=dice+roller`.

Condition	$\Delta\alpha$	ΔW	α_H	α_L	W_2	W_1	Sample
1	1	3 min	1	0	5 min	2 min	226
2	1	8 min	1	0	10 min	2 min	220
3	1	13 min	1	0	15 min	2 min	217
4	0.5	3 min	1	0.5	5 min	2 min	222
5	0.5	8 min	1	0.5	10 min	2 min	227
6	0.5	13 min	1	0.5	15 min	2 min	222
7	0.1	3 min	1	0.9	5 min	2 min	220
8	0.1	8 min	1	0.9	10 min	2 min	233
9	0.1	13 min	1	0.9	15 min	2 min	234

Table 1 Experimental conditions.

they must wait in a virtual queue, then answer a two-question survey in exchange for a 1 US dollar payment (the payment is independent of the wait time in queue). Participants were informed that the experiment could take up to 30 minutes. We exclude participants that did not complete the experiment. Moreover, while waiting in the virtual queue, and to ensure that participants experience the wait, they are asked to click a button that appears every 60 seconds in order to move ahead in the queue. The remaining waiting time for participants in a given queue stops from elapsing until they click that button. We recorded the time that participants took to click each button and excluded participants that, on average, took longer than 30 seconds to click those buttons once they appeared.

We set the target sample size for the experiment and our analysis plans a priori[6]. A total of 2,373 participants (44.33% female, mean age M_{age} = 37.97, standard deviation SD_{age} = 11.70) were recruited. In our experiment we have a completion rate of 95%. From our original sample, 90% of participants took, on average, less than 30 seconds to click the buttons, and the mean and median average click time was 14 seconds and 4 seconds, respectively. After exclusions, we are left

[6] For presentation and communication purposes, we slightly change the wording of the third pre-registered hypothesis, without modifying its content and scope. This has no effect on our conclusions. The pre-registration document can be found at: `https://aspredicted.org/blind.php?x=7u6ar3`.

with a sample of 2,021 participants (45.47% female, mean age $M_{age} = 38.23$, standard deviation $SD_{age} = 11.99$).

6.2. Hypotheses

The existence of intrinsic lying costs underpins our lying-aversion model. Indeed, without lying costs as a behavioural extension in the customer problem (2), our queueing game becomes trivial: We can clearly see that if $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) = 0$ for $y \neq x$, then all customers misreport with probability 1. In other words, the existence of lying costs is a necessary condition to retrieve any level of honesty. This leads to our first hypothesis:

H1. The proportion of participants who misreport their private information in order to wait in the Short queue is bounded away from 1.

As discussed before, motivated by the expressions for the misreporting probabilities under the sensible lying cost specifications in §5.2, we are interested in investigating the effects of $\Delta\alpha$ and ΔW. In relation to this, we note that the extant literature, which focuses on monetary payoffs, provides converging evidence that the average amount of lying does not change if the difference in monetary payoffs is increased from a few cents to 50 USD, a 500-fold increase (Abeler et al. 2019). That is, it presents consistent evidence that lying costs increase linearly in the monetary gain (Mazar et al. 2008). Indeed, according to our linear lying cost specification in (5), changes in both $\Delta\alpha$ and ΔW should not have an effect on misreporting. This leads to the following hypotheses:

H2. Changing the difference in the waiting times between the two queues does not have an effect on the proportion of participants who misreport.

H3. Changing the difference in the routing probabilities does not have an effect on the proportion of participants who misreport.

6.3. Experimental Results

We now describe our experimental results. In the FFH experimental design, since it is not possible to observe individual die outcomes, hypotheses regarding the misreporting behaviour of participants are tested by conducting statistical analyses on participants' reported outcomes; this is a standard procedure, see e.g., Abeler et al. (2019), Fischbacher and Föllmi-Heusi (2013). Indeed, under the assumption that there is no down-reporting (i.e., participants lying to their disadvantage), and since die outcomes are generated by a discrete uniform distribution in every experimental condition, differences in reporting behaviour between conditions can be attributed to differences in lying behaviour, and differences between the distribution of reports and the uniform distribution represent evidence for lying (Fries and Parra 2021, Fries et al. 2021). As a robustness check, in Appendix B.3.3 we conduct a simulation analysis where we investigate the effect of the sampling variation of the die outcome. Our simulation results show that for the sample sizes that we work with in the present experiment, we can confidently conclude that differences in reporting behaviour are indeed attributed to differences in lying behaviour, and not to sampling variation.

In Figure 1, we plot the proportions of participants who reported the number 5 across all experimental conditions. In Appendix B.1, we present corresponding proportions in Table 2. Recall that participants have the incentive to report the number 5 to reduce their waiting time in the queue.

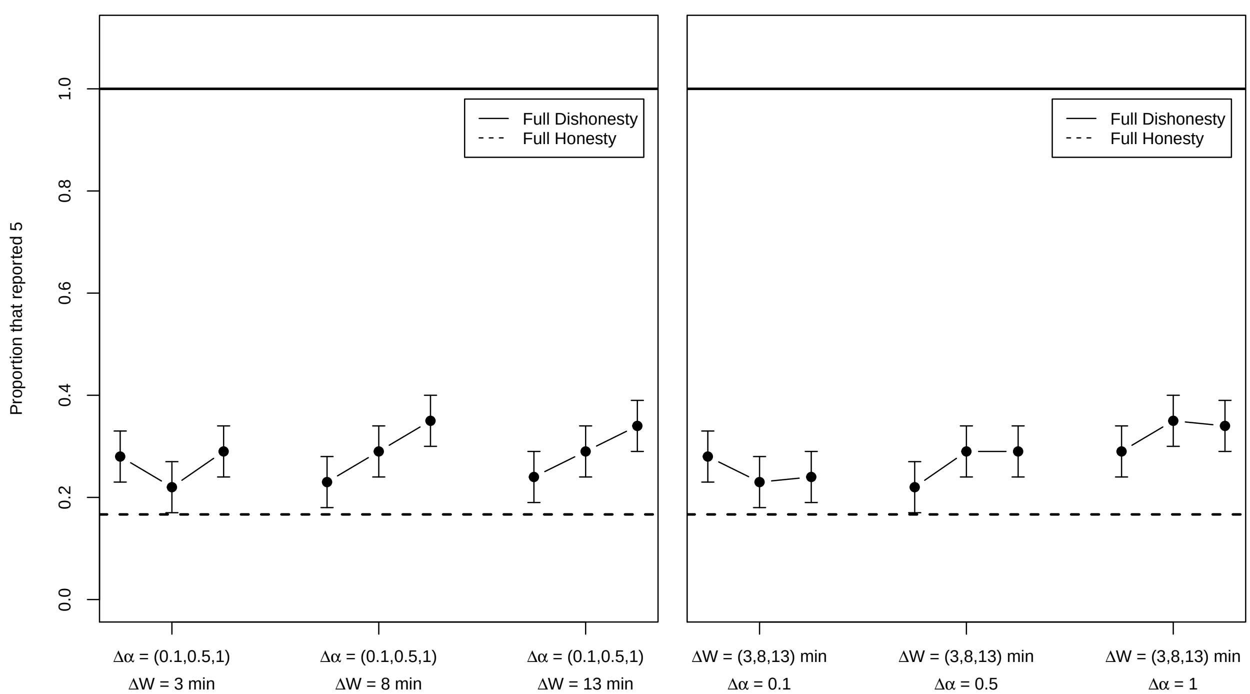

Figure 1 Proportions of participants who reported the number 5 across experimental conditions.

Human misreporting behaviour demonstrates the existence of lying costs. We run exact binomial tests for the proportion of participants who reported the number 5 in each condition compared to the proportion that would have reported the number 5 under full honesty i.e., 1/6, (p-values < 0.05). This means that participants misreported in all conditions, which can be readily seen in Figure 1, where the black horizontal dashed line represents the expected proportion of customers who would have reported the number 5 under the assumption of full honesty. Any value above the dashed line is based on untruthfulness, in expectation. Importantly, we find that participants misreport little across all experimental conditions. It is clear from Figure 1 that the proportion of participants who reported the number 5 is far from 1, i.e., the full-dishonesty black horizontal solid line in the figure. Since the FFH paradigm focuses on intrinsic costs and, by design, controls for other confounding costs, any reluctance to lie can be attributed to intrinsic lying costs. This result strongly supports the existence of lying costs, as per our model in (2), i.e., we find support for H1.

Impact of changes in waiting times. To test the effects of changes in waiting times, we run logistic regressions for the probability to claim the number 5, where we control for age and gender (see Table 3 in Appendix B.2 for details). We find that the coefficient for the difference in waiting times is not significant and close to 0 in all of the regression specifications. This provides strong evidence to support the claim that changes in waiting times affect very little (if at all) misreporting behaviour. That is, we find support for H2. In Appendix B.3, we present several non-parametric tests as robustness checks to further support our conclusions.

Impact of changes in the routing probabilities. In our logistic regressions, based on the Akaike Information Criterion (AIC), we find that the most appropriate model specification among the ones considered, is Model (2a) which includes an intercept, demographic covariates (insignificant in our data), and the $\Delta\alpha$ covariate (see Table 3 in Appendix B.2). We observe that the coefficient for $\Delta\alpha$ is significant and positive, indicating that the misreporting probability increases in $\Delta\alpha$. The other specifications included the ΔW covariate, and the interaction term $\Delta\alpha * \Delta W$, which represents the difference in waiting times between the priority classes $W_L - W_H$. Overall, we find sufficient statistical evidence in our data to support the claim that changes in $\Delta\alpha$ affect misreporting behaviour.

That is, we find support to reject H3. In Appendix B.3, we present several non-parametric tests as robustness checks to further support our conclusions.

6.4. Modelling Implications

The results in our experimental study indicate that the great majority of explainable lying behaviour can be attributed to $\Delta\alpha$, while only a very small portion (if at all) can be attributed to the expected waiting times ΔW. We note that none of the considered sensible models in §5.1 are consistent with both findings, i.e., their predictions for the best-response misreporting probabilities do not exhibit the same directional patterns as our experimental results. This indicates that we require to revise our assumptions to better capture the direction of changes in the observed decisions in our data. To do so, we highlight the fact that the considered lying-cost models above are based on the extant literature which has mainly focused on settings where the consequences of the reporting behaviour are certain. In contrast, in our queueing setting, due to the routing probabilities, the outcomes of both telling the truth and misreporting may be uncertain. Importantly, we highlight that for our experimental conditions 1-3, where outcomes are certain (since $\alpha_H = 1$ and $\alpha_L = 0$), the linear lying cost specification i.e., $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W}) = c_x(W_x - W_y)^+$, is the only one which is consistent with our experimental results, as it is the only one where changes in the waiting times do not affect misreporting. As previously discussed, this finding is consistent with the robust empirical results in literature that misreporting is not affected by changes in material incentives (Abeler et al. 2019).

We now propose a model which is consistent with our experimental data, i.e., which leads, simultaneously, to a misreporting probability which is insensitive to the waiting times yet sensitive to changes in the routing probabilities. To guarantee insensitivity to waiting times, it is sufficient to assume that the lying cost, $\ell(x, y, \boldsymbol{\alpha}, \mathbf{W})$, is proportional to the material gain from lying, i.e., proportional to $c_x(W_x - W_y)^+$. This assumption is commonly made in the literature (Mazar et al. 2008) to justify the insensitivity of the misreporting probability to changes in the material incentives. As such, we ensure that the lying cost balances out the incentive to misreport, so that indeed varying the difference in waiting times does not affect the probability of misreporting.

To capture the dependence of the misreporting probability on the routing parameters, we allow for the routing probabilities to affect the lying aversion, θ, directly. We note that this assumption is informed by the literature which argues that the lying aversion of individuals is malleable (Mazar et al. 2008, Rosenbaum et al. 2014, Abeler et al. 2014) and, importantly, that can be directly shaped by outcome uncertainty (Celse et al. 2019, Dugar et al. 2019, Steinel et al. 2022). For example, Celse et al. (2019) posits that for lower values of the probability to get a benefit by means of misreporting, individuals experience a lower lying aversion because e.g., lying is less likely to result in the obtainment of a benefit. This creates a sort of "psychological distance" from misreporting, which allows people to further relax their moral standards (Ariely 2012, Dugar et al. 2019).

Based on the above, we assume that a customer of true type $x \in \{H, L\}$ makes a claim $y \in \{H, L\}$ to minimize a total expected cost equal to the sum of the expected delay cost and the intrinsic lying cost:

$$\underset{y\in\{H,L\}}{Min}\ c_x W_y + \frac{\theta}{\tau(\boldsymbol{\alpha})} c_x (W_x - W_y)^+, \tag{7}$$

where the lying aversion θ is shaped by a function[7] $\tau(\boldsymbol{\alpha}) > 0$. Based on (7), the expected cost of a tagged x type customer with lying aversion θ' is $c_x W_x$, for an honest claim $y = x$, and $c_x W_y + \frac{\theta'}{\tau(\boldsymbol{\alpha})} c_x (W_x - W_y)^+$, for a dishonest claim $y \neq x$. That tagged customer selects the claim y which minimizes her expected cost. The tagged customer will never misreport if there is no incentive to do so (i.e., $W_x < W_y$ for $y \neq x$) since in this case we have $c_x W_x < c_x W_y + \frac{\theta'}{\tau(\boldsymbol{\alpha})} c_x (W_x - W_y)^+$. In contrast, if there is an incentive to misreport (i.e., $W_x > W_y$ for $y \neq x$), the tagged customer misreports whenever her lying aversion parameter θ' is sufficiently small:

$$\theta' \leq \frac{\tau(\boldsymbol{\alpha}) c_x (W_x - W_y)}{c_x (W_x - W_y)} = \tau(\boldsymbol{\alpha}) \quad \text{for} \quad y \neq x.$$

[7] We could model the effect of $\tau(\boldsymbol{\alpha})$ on θ as $\theta\tau(\boldsymbol{\alpha})$, or additively as $(\theta + \tau(\boldsymbol{\alpha}))$. Our choice in (7) is based on convenience since it yields a misreporting probability equal to $\Phi(\tau(\boldsymbol{\alpha}))$, such that we can impose assumptions on $\tau(\boldsymbol{\alpha})$ more naturally based on observed changes in the misreporting probability.

Since individual lying aversions, i.e., specific θ values, are not observable, it follows that a typical x type customer misreports with probability:

$$\mathbb{P}(Y=y|X=x)=\Phi\big(\tau(\boldsymbol{\alpha})\big) \quad \text{for} \quad y\neq x, \tag{8}$$

where we recall that Φ is the CDF of the lying aversion Θ. Note that the expression in (8) is consistent with our experimental results where we saw that participant lying behavior can be explained by the routing probabilities rather than the waiting times.

6.4.1. Properties of $\tau(\boldsymbol{\alpha})$. We make assumptions on $\tau(\boldsymbol{\alpha})$ to capture changes in the misreporting probability which are consistent with our experimental data. Let $\underline{\boldsymbol{\alpha}}=min\{\alpha_H,\alpha_L\}$, which represents, for those customers who have the incentive to misreport, the probability to gain priority by *means of truth-telling*. For example, in our experimental study, we have $\alpha_H>\alpha_L$ so that L type customers have the incentive to misreport. In this case, $\underline{\boldsymbol{\alpha}}=\alpha_L$ is the probability to be assigned to the short queue given an L claim. Based on this, we make the following assumptions:

Assumptions 1. We assume that $\tau(\boldsymbol{\alpha})$ in (7) and (8) satisfies $\tau(\cdot)$ is a function of $\underline{\boldsymbol{\alpha}}$ alone, $\tau'(\underline{\boldsymbol{\alpha}})\leq 0$, $\tau''(\underline{\boldsymbol{\alpha}})\geq 0$, and $\tau(0)<\bar{\theta}$.

The first assumption, that $\tau(\cdot)$ is a function of $\underline{\boldsymbol{\alpha}}$ only, means that lying aversion is shaped only by the probability to get priority by means of truth-telling. We note that in our experimental study we kept α_H fixed, so the observed effect can only be attributed to changes in α_L (see Table 4 in Appendix B.2 for details). In Appendix C, we present an additional experiment where we vary both routing probabilities. While the misreporting probability does decrease in the probability to get priority by means of misreporting (Celse et al. 2019), i.e., α_H in the experiment, we find that this effect is not significant (see Table 10 in Appendix C.1 for details). Thus, for simplicity and as a conservative assumption[8], we assume that the lying aversion is shaped only by the probability to get priority by means of truth-telling, i.e., α_L in the experiment.

[8] We will see that, in the optimal solution, the Manager sets $\alpha_H^*=1$. If we also allow for the misreporting probability to decrease in α_H, as indicated by the coefficients in Table 10 in Appendix C.1, and consistently with the hypothesized effect in Celse et al. (2019), then we will continue to obtain that $\alpha_H^*=1$.

The second assumption, $\tau'(\underline{\boldsymbol{\alpha}}) \leq 0$, implies that the lying aversion increases in $\underline{\boldsymbol{\alpha}}$, which is consistent with our experimental observation that misreporting decreases in $\underline{\boldsymbol{\alpha}} = \alpha_L$ (see Table 4 in Appendix B.2 for details). This can be interpreted as follows. For higher values of $\underline{\boldsymbol{\alpha}}$, individuals experience a higher lying aversion because e.g., there is a higher opportunity to obtain priority and report truthfully at the same time, so they feel worse when they lie because they know that they could have obtained priority by honest means instead. Indeed, individuals' intrinsic preferences for truth-telling are well recognized (Abeler et al. 2019). Note that in this assumption we have a weak inequality. The reason is that we argue that the sensitivity to $\underline{\boldsymbol{\alpha}}$ (i.e., the strength of the effect) may depend on the population under analysis. Indeed, in our additional experiment in Appendix C, we observe a difference in the honesty attitudes (see Figure 2 in Appendix C.1): This population was more honest in comparison to the one presented in §6. While we observe negative coefficients for $\underline{\boldsymbol{\alpha}}$ (see Table 4 in Appendix B.2 for details), which is consistent with the experimental findings in §6, the effect was not strong enough to be significant. For generality, we allow such weak inequality for our results to be contingent on the responsiveness to $\underline{\boldsymbol{\alpha}}$ of a particular population (we will formalize this concept later in §7.3).

The third assumption, $\tau''(\underline{\boldsymbol{\alpha}}) \geq 0$, captures a constant/diminishing marginal effect of $\underline{\boldsymbol{\alpha}}$ on misreporting. In particular, this assumption is made for analytical tractability as it leads to the unimodality in the waiting cost. This assumption is not restrictive since it does not change the main insights derived from the identified optimal scheduling policy (see Appendix A.2 for details).

Finally, the last assumption, $\tau(0) < \bar{\theta}$, ensures that not everyone misreports, which is consistent with our experimental results, the extant literature, and our motivating examples. To see this, recall that $\bar{\theta}$ represents the upper bound of the support for the lying aversion distribution. Since $\tau(\underline{\boldsymbol{\alpha}})$ decreases in $\underline{\boldsymbol{\alpha}}$, if $\tau(0) < \bar{\theta}$ then it follows that $\Phi(\tau(\underline{\boldsymbol{\alpha}})) < 1$, that is, that not everyone misreports.

Overall, our lying aversion model in (7) captures the fact that while customers have a higher incentive to lie when the difference in waiting times between the queues is increased, they also feel worse about lying because the stakes are higher. Thus, variations in waiting times do not affect misreporting. At the same time, due to $\underline{\boldsymbol{\alpha}}$, changes in the probability to get priority by means of truth-telling affect customer misreporting behaviour as it shapes customers' lying aversion directly according to $\tau(\underline{\boldsymbol{\alpha}})$. Finally, in Appendix B.4, we conduct a structural estimation analysis where we illustrate how a simple specification for $\tau(\boldsymbol{\alpha})$ as per our Assumptions 1 can achieve a better fit to the experimental data in comparison to the described sensible lying costs specifications.

7. Optimal Scheduling Policy

In this section, we derive the optimal routing probabilities, α_H^* and α_L^*, which minimize the expected waiting cost in the system, at equilibrium.

7.1. Over and Under-Prioritization Trade-Off

We begin by elaborating on the main trade-off that the Manager faces, to build intuition. In the Manager's problem (1), there is an underlying cost asymmetry between the cost that arises from under-prioritization (i.e., H type customers who are placed in the regular queue), and from over-prioritization (i.e., L type customers who are placed in the priority queue). To see this, we let $\delta_H = \mathbb{P}(K=2|X=H)$ be the steady-state under-prioritization probability and $\delta_L = \mathbb{P}(K=1|X=L)$ be the steady-state over-prioritization probability, given by:

$$\delta_H = (1-\alpha_H)\mathbb{P}(Y=H|X=H) + (1-\alpha_L)\mathbb{P}(Y=L|X=H), \tag{9}$$

$$\delta_L = \alpha_H \mathbb{P}(Y=H|X=L) + \alpha_L \mathbb{P}(Y=L|X=L). \tag{10}$$

For any given customer reporting behaviour $\mathbb{P}(Y=y|X=x)$, we study, in the following lemma, the monotonicity of the expected waiting cost as a function of those prioritization errors.

Lemma 1. *The expected waiting cost in the system C_s in (1) is increasing in both δ_H and δ_L. Moreover, it is more sensitive to δ_H than δ_L, i.e., $\frac{\partial C_s}{\partial \delta_H} > \frac{\partial C_s}{\partial \delta_L} > 0$ if, and only if, $\delta_H + \delta_L < 1$. The additional sensitivity to δ_H over δ_L, i.e., $\frac{\partial C_s}{\partial \delta_H} - \frac{\partial C_s}{\partial \delta_L}$, increases in ρ.*

Lemma (1) is similar to Lemma 7 in Singh et al. (2021) and Proposition 2 in Argon and Ziya (2009). We repeat this result here to draw parallels between our setting, where customers are strategic in their claims, and the models in those papers where customers do not have the ability to influence their priority classifications. Consistent with intuition, Lemma 1 shows that the waiting cost increases strictly in the prioritization errors δ_H and δ_L under a condition which, we will see, is satisfied at optimum; see Proposition 1 in the next section. Importantly, Lemma 1 demonstrates the asymmetric impact of under-prioritization versus over-prioritization errors, particularly in congested systems: Under-prioritization is more costly than over-prioritization. Given this asymmetry in the Manager's trade-off, we will see that the optimal policy allows to eliminate the over-prioritization error, while minimizing the over-prioritization error.

7.2. Game Equilibrium

Recall that the probability density function of the lying aversion Θ is denoted by ϕ, and its cumulative distribution function by Φ. The failure rate of Θ is given by $h(z) = \phi(z)/(1-\Phi(z))$. We assume that the lying aversion distribution is an *increasing failure rate* (IFR) distribution. The random variable Θ is said to have an increasing failure rate if $h(z)$ is weakly increasing in z for $\Phi(z) < 1$; see Lariviere and Porteus (2001). We note that the assumption that the lying aversion Θ is IFR is not restrictive because it captures many common distributions, e.g., uniform, exponential, normal, among others (Banciu and Mirchandani 2013). Also, under the IFR assumption, the Manager's objective function is unimodal and there is a unique equilibrium in the game.

Proposition 1. *Based on problems (1) and (7), under Assumptions 1, if the lying aversion Θ is IFR, then there exists a unique equilibrium which arises in the game between the customers and the Manager:*

(a) All customers with true type $X = H$ make a truthful claim $Y = H$.

(b) A fraction $\Phi(\tau(\alpha_L^*)) < 1$ of customers with true type $X = L$ are dishonest i.e., claim $Y = H$.

(c) The optimal routing policy is to assign high priority to all high-claim customers, i.e., $\alpha_H^* = 1$, and to assign high priority to low-claim customers with probability α_L^*. We have that $\alpha_L^* = 0$ if, and only if, $1 + \tau'(0)h(\tau(0)) \geq 0$, and $\alpha_L^* \in (0,1)$ as the unique solution to $1 + (1 - \alpha_L^*)\tau'(\alpha_L^*)h(\tau(\alpha_L^*)) = 0$, otherwise.

Consistent with intuition, part (a) of Proposition 1 shows that customers of type $X = H$ always claim their true type. Part (b) shows that only a proportion (bounded away from 1) of $X = L$ type customers misreport. This result is consistent with the extant literature and with our motivating examples. Finally, part (c) shows that the optimal policy is not FCFS (which corresponds to $\alpha_H = \alpha_L = 1$). Indeed, as long as there is some level of honesty in the system, the Manager is able to extract, at equilibrium, some information from customer claims about the true customer types:

$$\mathbb{P}^*(X = H | Y = H) = \frac{p_H}{p_H + p_L \Phi(\tau(\alpha_L^*))} > p_H,$$
$$\mathbb{P}^*(X = L | Y = L) = 1 > p_L,$$

where we recall that p_H and p_L are the proportions of true H and L types. This shows that due to customers' lying aversion, the Manager can still extract some information from the customers' own claims, despite the prevalence of misreporting. This is a useful insight for managers. For example, in one of our motivating examples, the National Health Service, anticipating dishonest claims from patients, dropped a chat-bot symptom checker application designed to aid in booking appointments (Heather 2017). Our theoretical results provide insight into when this may, or may not, be a good decision.

In Appendix D, we derive general conditions for the misreporting probability, as a function of both routing probabilities, under which our optimal scheduling policy results would continue to hold. This provides sufficient conditions for which our results hold under more general specifications of $\tau(\boldsymbol{\alpha})$ or under different formalizations for the customer problem.

7.3. Honour System and Upgrading Policy

From Proposition 1, we can see that $\mathbb{P}^*(Y = H | X = H) = 1$ and $\mathbb{P}^*(Y = H | X = L) = \Phi(\tau(\alpha_L)) < 1$. Thus, at equilibrium, from (9) and (10), we can see that the Manager faces the following prioritization errors:

$$\delta_H = 1 - \alpha_H, \tag{11}$$

$$\delta_L = \alpha_H \Phi(\tau(\alpha_L)) + \alpha_L (1 - \Phi(\tau(\alpha_L))). \tag{12}$$

In part (c) of Proposition 1, we see that, on one hand, the system Manager should always set $\alpha_H^* = 1$, i.e., to *eliminate* the under-prioritization error δ_H; see (11), even if this allows more misreporting customers to jump the queue; see (12). This is driven by the asymmetry between the costs of under-prioritization and over-prioritization in Lemma 1. This shows that Managers should exert caution when adopting strategies to mitigate misreporting at the expense of under-prioritizing. For example, this was the case in the groceries delivery services example of the introduction, where it took too long to verify the vulnerability status of people in order to give them priority slots to gain access to food (Syal 2021).

On the other hand, in part (c) of Proposition 1, we see that, the Manager uses $\alpha_L^* \geq 0$ in order to *minimize* the over-prioritization error δ_L. From (12), we find two effects for increasing α_L. First, increasing α_L leads to an increase in the second part of δ_L due to a higher proportion of honest L customers $(1-\Phi(\tau(\alpha_L)))$ being *upgraded* to the high-priority queue. Second, increasing α_L leads to a decrease in the first part of δ_L since it allows to *incentivize more honesty*: A lower proportion of L customers $\Phi(\tau(\alpha_L))$ misreport due to their lying aversion. Therefore, to minimize the over-prioritization error δ_L, the Manager should balance that tension optimally by deciding *how much honesty to incentivize* from customers.

To solve for the optimal level of honesty in the system, we find that a particular metric of semi-elasticity, i.e., a measure of the sensitivity of one variable to another, arises naturally. The semi-elasticity of w with respect to v, $\mathcal{S}_w(v) = \frac{dw/dv}{w}$, measures the percentage change in w in terms of a change (not percentage-wise) in v (Wooldridge 2015). If we let $\xi = 1-\Phi(\tau(\alpha_L))$ be the proportion of honest customers, then we have that the *semi-elasticity of honest behaviour* with respect to changes in the upgrading probability α_L is given by:

$$\mathcal{S}_\xi(\alpha_L) = -\frac{\tau'(\alpha_L)\phi(\tau(\alpha_L))}{1-\Phi(\tau(\alpha_L))} = -\tau'(\alpha_L)h(\tau(\alpha_L)) \geq 0. \tag{13}$$

It is easy to see that $\mathcal{S}_\xi(\alpha_L)$ weakly decreases in the upgrading probability α_L, which implies that the maximum semi-elasticity of honest behaviour is given by $\mathcal{S}_\xi(0)$. In part (c) of Proposition 1, we see that upgrading is optimal whenever $1+\tau'(0)h(f(0)) < 0$, which is equivalent to $\mathcal{S}_\xi(0) > 1$.

Intuitively, the condition $\mathcal{S}_\xi(0) > 1$ guarantees that there is some α_L for which the semi-elasticity of honest behaviour is sufficiently strong so that upgrading decreases the overall cost.

Finally, if the population under analysis does not respond to the upgrading control sufficiently, or does not respond at all (i.e., $\tau'(\alpha_L) = 0$), Proposition 1 prescribes an honour system, where priorities are given according to customer claims. For example, for the population under analysis in our additional experimental investigation in Appendix C, it is fair to say that $\tau'(\alpha_L) = 0$. Based on this, Proposition 1 would prescribe an honour system for such a population. This provides support to the honour-based scheduling policies observed in practice. For example, the UK used an online booking system to get access to COVID-19 tests at home, which relied on public honesty in claiming their employment status (Weaver and Proctor 2020).

7.4. Upgrading as a Deviation from the cμ Rule

It is well known that, when customer types are fully known, that the celebrated $c\mu$ rule minimizes the expected cost in the system. In our game, the true priority level of an arriving customer (equivalently, the true index $c\mu$) is unknown to the system Manager. Based on customer claims, in equilibrium, the system Manager can infer the expected $c\mu$ customer index, $I(y) = c_H\mu\mathbb{P}(X = H|Y = y) + c_L\mu\mathbb{P}(X = L|Y = y)$. It is easy to see that $I(H) > I(L)$ if, and only if, $c_H > c_L$, which is our assumption throughout this paper. When the uncertainty in customer types is exogenous, Argon and Ziya (2009) argue that customers should be prioritized in decreasing order of the expected $c\mu$ indices. In our model, upgrading is optimal because customer claims are *endogenous*: By allowing for some L claim customers to be in the high priority lane (i.e., by deviating from the $c\mu$ rule), the Manager incites customers to be more honest in their reporting, which ultimately benefits the system as a whole.

8. Conclusions

We studied the extent to which people misreport their private information in a two-class priority queueing system, where priorities are assigned according to customer claims, and where these claims are not verifiable nor punishable. Although standard economic theory predicts that all customers will claim that they should be given high priority, we showed, through the analysis of a theoretical queueing model and controlled experiments, that human behaviour deviates far from this prediction. We theorized and experimentally validated the intrinsic costs that people incur when they lie. We studied the consequences of uncertainty (probabilistic routing) on lying, and found that the scheduling policy, rather than the waiting times themselves, significantly impacts lying behaviour.

Priority systems which rely on customer claims are prevalent in practice. Our results highlight the informative value of those claims, even in systems where customers can lie with total impunity. Thus, we provided theoretical and experimental evidence that customer claims should be sought by the system Manager. This is particularly relevant to priority settings where market mechanisms, e.g., pricing the relevant services, are not applicable. In particular, we found theoretical evidence that supports the current honour scheduling system (routing customers according to their claims) which is in place in many real settings. Additionally, we found that, when the lying aversion cost is sufficiently elastic, a different prioritization rule, which deviates from the celebrated $c\mu$ rule, is optimal. Essentially, this rule prescribes upgrading some low-priority claims to the high-priority class in order to incentivize honesty in the system as a whole. We highlight that there is no guarantee, with other scheduling-policy implementations, to be able to identify operational controls that incentivize more honesty. One contribution of our work is, therefore, that we both experimentally and theoretically show that the considered α-policy provides a useful operational control (i.e., probabilistic routing) to mitigate dishonesty.

Limitations and future research. The study of untruthfulness in queues is rich and strongly motivated by practice. We hope that our work will spark more interest in this domain. Our experimental results are based on simple two-priority queueing settings. We advocate conducting more experimental work to further understand the role of intrinsic preferences for truth-telling in the reporting behaviour of people in more complicated queueing settings. Future work can, for example, investigate the extent to which people report untruthfully in a multi-priority queueing system, and unveil the relevant determinants in that case. While we foresee that some people will be honest and others will be untruthful, it is not clear, a priori, how customer claims will be distributed across the multiple priority queues: For example, will there be partial lying in this case?

In this paper, we focused on priority settings where the dynamic state of the queue is unobservable to customers, i.e., the length of the queue and the real-time behaviour of other customers. We also assumed the absence of reputation concerns since the service setting involves an anonymous, one-shot, type of interaction with the Manager. These properties are common features shared by many priority queueing settings, but certainly not all systems. Future research could focus on understanding customer untruthfulness in priority queues where the dynamic state of the queue is observable, e.g., in emergency departments where patients may exaggerate their own symptoms at the triage stage to receive medical service faster. In such settings, where customers observe each other, reputation and negative externality concerns may play a major role. Also, as mentioned above, this work is limited to settings where customers have a rare or one-shot interaction with the service system, like the ones in our motivating examples. We advocate that future research studies the effect of repeated interactions, where we envision that other behavioural factors may play an important role in the misreporting behaviour.

Finally, in this work, we restricted attention to the α-policy and found the optimal way to schedule customers based on their claims. More generally, different policy implementations, that tap into different behavioural factors, may be considered. The study of lying behaviour under different scheduling policies is an interesting future research direction.

References

Abeler, J., A. Becker, and A. Falk (2014). Representative evidence on lying costs. *Journal of Public Economics 113*, 96–104.

Abeler, J., D. Nosenzo, and C. Raymond (2019). Preferences for truth-telling. *Econometrica 87*(4), 1115–1153.

Afeche, P. (2013). Incentive-compatible revenue management in queueing systems: Optimal strategic delay. *Manufacturing & Service Operations Management 15*(3), 423–443.

Afeche, P. and H. Mendelson (2004). Pricing and priority auctions in queueing systems with a generalized delay cost structure. *Management science 50*(7), 869–882.

Afeche, P. and J. M. Pavlin (2016). Optimal price/lead-time menus for queues with customer choice: Segmentation, pooling, and strategic delay. *Management Science 62*(8), 2412–2436.

Agresti, A. (2003). *Categorical data analysis*, Volume 482. John Wiley & Sons.

Allon, G. and E. Hanany (2012). Cutting in line: Social norms in queues. *Management Science 58*(3), 493–506.

Allon, G. and M. Kremer (2018). Behavioral foundations of queueing systems. *The Handbook of Behavioral Operations 9325*, 325–366.

Althenayyan, A., S. Cui, S. Ulku, and L. Yang (2022). Not all lines are skipped equally: An experimental investigation of line-sitting and express lines. *Available at SSRN 4179751*.

Argon, N. T. and S. Ziya (2009). Priority assignment under imperfect information on customer type identities. *Manufacturing & Service Operations Management 11*(4), 674–693.

Ariely, D. (2012). *The (honest) truth about dishonesty: : how we lie to everyone–especially ourselves.* New York: Harper Collins Publishers.

Armony, M., G. Roels, and H. Song (2021). Pooling queues with strategic servers: The effects of customer ownership. *Operations Research 69*(1), 13–29.

Banciu, M. and P. Mirchandani (2013). New results concerning probability distributions with increasing generalized failure rates. *Operations Research 61*(4), 925–931.

BBC (2020a). Coronavirus: People without symptoms 'misusing testing'. https://www.bbc.co.uk/news/health-54088206.

BBC (2020b). Matt hancock has urged coronavirus tests to be prioritised for key workers. jade says she lied about her employment status to get a test. https://twitter.com/BBCBreakfast/status/1306109442295988224.

Bren, A. and S. Saghafian (2019). Data-driven percentile optimization for multiclass queueing systems with model ambiguity: Theory and application. *INFORMS Journal on Optimization 1*(4), 267–287.

Buell, R. W. (2021). Last-place aversion in queues. *Management Science 67*, 1430–1452.

Celse, J., S. Max, W. Steinel, I. Soraperra, and S. Shalvi (2019). Uncertain lies: How payoff uncertainty affects dishonesty. *Journal of Economic Psychology 71*, 117–125.

Coffman Jr, E. G. and I. Mitrani (1980). A characterization of waiting time performance realizable by single-server queues. *Operations Research 28*(3-part-ii), 810–821.

Cox, D. R. and W. Smith (1991). *Queues*, Volume 2. CRC Press.

Dai, Z., F. Galeotti, and M. C. Villeval (2018). Cheating in the lab predicts fraud in the field: An experiment in public transportation. *Management Science 64*(3), 1081–1100.

DellaVigna, S., J. A. List, U. Malmendier, and G. Rao (2016). Voting to tell others. *The Review of Economic Studies 84*(1), 143–181.

Duch, R. M., D. Laroze, and A. Zakharov (2021). The moral cost of lying. Technical report, University of Oxford.

Dufwenberg, M. and M. A. Dufwenberg (2018). Lies in disguise–a theoretical analysis of cheating. *Journal of Economic Theory 175*, 248–264.

Dugar, S., A. Mitra, and Q. Shahriar (2019). Deception: The role of uncertain consequences. *European Economic Review 114*, 1–18.

Fischbacher, U. and F. Föllmi-Heusi (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association 11*(3), 525–547.

Fries, T., U. Gneezy, A. Kajackaite, and D. Parra (2021). Observability and lying. *Journal of Economic Behavior & Organization 189*, 132–149.

Fries, T. and D. Parra (2021). Because i (don't) deserve it: Entitlement and lying behavior. *Journal of Economic Behavior & Organization 185*, 495–512.

Gächter, S. and J. F. Schulz (2016). Intrinsic honesty and the prevalence of rule violations across societies. *Nature 531*(7595), 496–499.

Gibson, R., C. Tanner, and A. F. Wagner (2013). Preferences for truthfulness: Heterogeneity among and within individuals. *American Economic Review 103*(1), 532–48.

Gneezy, U., A. Kajackaite, and J. Sobel (2018). Lying aversion and the size of the lie. *American Economic Review 108*(2), 419–53.

Gneezy, U., B. Rockenbach, and M. Serra-Garcia (2013). Measuring lying aversion. *Journal of Economic Behavior & Organization 93*, 293–300.

Hassin, R. and M. Haviv (2003). *To queue or not to queue: Equilibrium behavior in queueing systems*, Volume 59. Springer Science & Business Media.

Heather, B. (2017). Babylon pilot was not pursued by ccgs. `https://www.hsj.co.uk/technology-and-innovation/babylon-pilot-was-not-pursued-by-ccgs/7021131.article`.

Hu, M., R. Momot, and J. Wang (2021). Privacy management in service systems. Technical report, University of Toronto.

Hugh-Jones, D. (2019). True lies. *Journal of the Economic Science Association 5*(2), 255–268.

Jones, J. (2020). Take it from a paramedic: the ambulance service can't keep up with demand. `https://www.theguardian.com/commentisfree/2020/jan/29/paramedic-ambulance-service-call-triage`.

Kaidan, M. (2020). Coronavirus uk: Supermarkets to prioritise vulnerable people with deliveries. https://www.express.co.uk/news/uk/1261014/coronavirus-supermarket-news-vulnerable-delivery-priority-shopping-sainsburys-waitrose.

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies 76*(4), 1359–1395.

Khalmetski, K. and D. Sliwka (2019). Disguising lies—image concerns and partial lying in cheating games. *American Economic Journal: Microeconomics 11*(4), 79–110.

Kim, S.-H., J. Tong, and C. Peden (2020). Admission control biases in hospital unit capacity management: How occupancy information hurdles and decision noise impact utilization. *Management Science 66*(11), 5151–5170.

Kirton, J. A., W. Thompson, M. Pearce, and J. M. Brown (2020). Ability of the wider dental team to triage patients with acute conditions: a qualitative study. *British dental journal 228*(2), 103–107.

Kittsteiner, T. and B. Moldovanu (2005). Priority auctions and queue disciplines that depend on processing time. *Management Science 51*(2), 236–248.

Kleinrock, L. (1967). Optimum bribing for queue position. *Operations Research 15*(2), 304–318.

Lariviere, M. A. and E. L. Porteus (2001). Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing & service operations management 3*(4), 293–305.

Lui, F. T. (1985). An equilibrium queuing model of bribery. *Journal of political economy 93*(4), 760–781.

Luo, J., L. Valdés, and S. Linardi (2022). Experienced and prospective wait in queues: A behavioral investigation. *Available at SSRN 4169028*.

Mazar, N., O. Amir, and D. Ariely (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of marketing research 45*(6), 633–644.

Mendelson, H. and S. Whang (1990). Optimal incentive-compatible priority pricing for the m/m/1 queue. *Operations research 38*(5), 870–883.

NPR (2021a). How people are jumping the covid-19 vaccine line. `https://www.npr.org/sections/health-shots/2021/02/09/965841419/does-loose-enforcement-of-vaccine-eligibility-rules-encourage-line-jumping`.

NPR (2021b). Is it ever ok to jump ahead in the vaccine line? `https://www.npr.org/sections/health-shots/2021/02/06/964139633/is-it-ever-ok-to-jump-ahead-in-the-vaccine-line?t=1623992132881`.

Özer, Ö., Y. Zheng, and K.-Y. Chen (2011). Trust in forecast information sharing. *Management Science 57*(6), 1111–1137.

Rosenbaum, S. M., S. Billinger, and N. Stieglitz (2014). Let's be honest: A review of experimental evidence of honesty and truth-telling. *Journal of Economic Psychology 45*, 181–196.

Shunko, M., J. Niederhoff, and Y. Rosokha (2018). Humans are not machines: The behavioral impact of queueing design on service time. *Management Science 64*(1), 453–473.

Singh, S., I. Gurvich, and J. Van Mieghem (2021). Feature-based design of priority queues: Digital triage in healthcare. Technical report, Northwestern University.

Steinel, W., K. Valtcheva, J. Gross, J. Celse, S. Max, and S. Shalvi (2022). (dis) honesty in the face of uncertain gains or losses. *Journal of Economic Psychology 90*, 102487.

Syal, R. (2021). Thousands of vulnerable people missed out on support at start of pandemic. https://www.theguardian.com/world/2021/apr/21/thousands-of-vulnerable-people-missed-out-on-support-at-start-of-pandemic.

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge university press.

Ülkü, S., C. Hydock, and S. Cui (2020). Making the wait worthwhile: Experiments on the effect of queueing on consumption. *Management Science 66*(3), 1149–1171.

Van der Zee, S. and H. Theil (1961). Priority assignment in waiting-line problems under conditions of misclassification. *Operations Research 9*(6), 875–885.

Van Mieghem, J. A. (1995). Dynamic scheduling with convex delay costs: The generalized c— mu rule. *The Annals of Applied Probability 5*(3), 809–833.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: Journal of the Econometric Society*, 307–333.

Wang, J. and Y.-P. Zhou (2018). Impact of queue configuration on service time: Evidence from a supermarket. *Management Science 64*(7), 3055–3075.

Weaver, M. and K. Proctor (2020). No checks to be made on essential-worker status for uk covid-19 tests. `https://www.theguardian.com/world/2020/apr/24/no-checks-to-be-made-on-essential-worker-status-for-uk-covid-19-tests`.

Wooldridge, J. M. (2015). *Introductory econometrics: A modern approach.* Cengage learning.

Yang, L. (2021). Invite your friend and you'll move up in line: optimal design of referral priority programs. *Manufacturing & Service Operations Management 23*(5), 1139–1156.

Yang, L., S. Cui, and Z. Wang (2021). Design of covid-19 testing queues. *Production and Operations Management*.

Appendix A: Technical Proofs

A.1. Lemma 1

PROOF. Recall that, by assumption, we have that $c_H > c_L$. We can write the expected waiting cost in the system (1) as:

$$C_s(\delta_H, \delta_L) = g(\rho) f(\delta_H, \delta_L),$$

where

$$g(\rho) = \left(\frac{\rho^2}{1-\rho}\right),$$

$$f(\delta_H, \delta_L) = \frac{p_H c_H (1-\rho(1-\delta_H)) + p_L c_L (1-\rho\delta_L)}{p_H(1-\rho(1-\delta_H)) + p_L(1-\rho\delta_L)}.$$

After some algebraic manipulations we find:

$$\frac{\partial f}{\partial \delta_H} = \frac{\rho p_H p_L (c_H - c_L)(1-\rho\delta_L)}{(p_H(1-\rho(1-\delta_H)) + p_L(1-\rho\delta_L))^2},$$

$$\frac{\partial f}{\partial \delta_L} = \frac{\rho p_H p_L (c_H - c_L)(1-\rho(1-\delta_H))}{(p_H(1-\rho(1-\delta_H)) + p_L(1-\rho\delta_L))^2}.$$

Since $\rho \in (0,1)$ and $\delta_L, \delta_H \in [0,1]$, it follows that $\frac{\partial C_s}{\partial \delta_H} = g(\rho)\frac{\partial f}{\partial \delta_H} > 0$ and $\frac{\partial C_s}{\partial \delta_L} = g(\rho)\frac{\partial f}{\partial \delta_L} > 0$.

We define $\Delta f' \doteq \frac{\partial f}{\partial \delta_H} - \frac{\partial f}{\partial \delta_L}$. After some algebraic manipulations we find:

$$\Delta f' = \frac{p_H p_L (c_H - c_L)\rho^2 (1-\delta_H-\delta_L)}{(p_H(1-\rho(1-\delta_H)) + p_L(1-\rho\delta_L))^2}.$$

Since $\rho \in (0,1)$ we have that $\Delta C_s' = \frac{\partial C_s}{\partial \delta_H} - \frac{\partial C_s}{\partial \delta_L} = g(\rho)\Delta f' > 0$ if and only if $\delta_H + \delta_L < 1$.

Finally, taking the following derivatives:

$$\frac{\partial g(\rho)}{\partial \rho} = \frac{2-\rho}{(1-\rho)\rho} g(\rho),$$

$$\frac{\partial \Delta f'}{\partial \rho} = \frac{2 p_H p_L (c_H - c_L)\rho (1-\delta_H-\delta_L)}{(p_H(1-\rho(1-\delta_H)) + p_L(1-\rho\delta_L))^3},$$

and recalling that $\Delta C_s' = \frac{\partial C_s}{\partial \delta_H} - \frac{\partial C_s}{\partial \delta_L} = g(\rho)\Delta f'$, we have that

$$\frac{\partial \Delta C_s'}{\partial \rho} = \frac{\partial g(\rho)}{\partial \rho}\Delta f' + g(\rho)\frac{\partial \Delta f'}{\partial \rho}.$$

Since $\rho \in (0,1)$, we have that $g(\rho) > 0$ and $\frac{\partial g(\rho)}{\partial \rho} > 0$. Moreover, notice that the sign of both $\frac{\partial \Delta f'}{\partial \rho}, \Delta f$ is dictated by the sign of $1-\delta_H-\delta_L$. It follows that $\frac{\partial \Delta C_s'}{\partial \rho} > 0$ if and only if $\frac{\partial \Delta f'}{\partial \rho}, \Delta f' > 0 \iff \delta_H + \delta_L < 1$. ■

A.2. Proposition 1

Proof. There are 3 different cases to analyse: (1) when $\alpha_H < \alpha_L$, (2) when $\alpha_H = \alpha_L$, and (3) when $\alpha_H > \alpha_L$. Based on customers problem (7) it is easy to see that in the first case since $W_L < W_H$, all L type customers report their type and a proportion $\Phi(\tau(\alpha_H)) < 1$ of H type customers misreport (note that here α_H represents the probability to get a benefit through honest means). For any $\alpha_H < \alpha_L$, L type customers are given priority over H type customers, recall that $c_H > c_L$ thus the waiting cost in the system is higher than in a FCFS policy. Also, for the second case since $W_L = W_H$, irrespective of customer claiming behaviour, the waiting cost in the system is equal to a FCFS policy. We now investigate the third case in detail.

For $\alpha_H > \alpha_L$ we have that all H type customers report their type and a proportion $\Phi(\tau(\alpha_L)) < 1$ of L type customers misreport in equilibrium. This comes from the fact that the misreporting probability (8) does not depend on waiting times (otherwise the customer equilibrium would come from the solution of a fixed point problem). This result is intuitive since customers can affect each other only through the waiting times. Since waiting times do not affect customer misreporting behaviour, customers make their decisions irrespective of how others behave. Based on this, the Manager correctly anticipates the following prioritization error probabilities:

$$\delta_H(\alpha_H,\alpha_L) = (1-\alpha_H)\mathbb{P}(Y=H|X=H) + (1-\alpha_L)\mathbb{P}(Y=L|X=H) = 1-\alpha_H,$$

$$\delta_L(\alpha_H,\alpha_L) = \alpha_H\mathbb{P}(Y=H|X=L) + \alpha_L\mathbb{P}(Y=L|X=L) = \alpha_L + (\alpha_H-\alpha_L)\Phi(\tau(\alpha_L)).$$

In this proof, for any function $q(\cdot)$, we use $q'_z(\cdot)$ and $q''_z(\cdot)$ to denote, respectively, the first and the second partial derivative of $q(\cdot)$ with respect to z. The Manager defines the routing probabilities $\alpha_H, \alpha_L \in [0,1]$ in order to minimize the expected waiting cost C_s in the system which can be written as:

$$C_s = c\cdot f(\alpha_H,\alpha_L),$$

where

$$c = \left(\frac{\rho^2}{1-\rho}\right) > 0,$$
$$f(\alpha_H,\alpha_L) = \frac{p_H c_H(1-\rho(1-\delta_H(\alpha_H,\alpha_L))) + p_L c_L(1-\rho\delta_L(\alpha_H,\alpha_L))}{p_H(1-\rho(1-\delta_H(\alpha_H,\alpha_L))) + p_L(1-\rho\delta_L(\alpha_H,\alpha_L))},$$
$$f'_{\alpha_H}(\alpha_H,\alpha_L) = \frac{d(\alpha_H,\alpha_L)\gamma(\alpha_H,\alpha_L)}{g(\alpha_H,\alpha_L)},$$
$$f'_{\alpha_L}(\alpha_H,\alpha_L) = \frac{d(\alpha_H,\alpha_L)(1-\rho\alpha_H)(1-\Phi(\tau(\alpha_L)))\kappa(\alpha_H,\alpha_L)}{g(\alpha_H,\alpha_L)},$$

$$d(\alpha_H,\alpha_L)=p_Hp_L\rho(c_H-c_L)>0,$$

$$g(\alpha_H,\alpha_L)=(p_H(1-\rho(1-\delta_H(\alpha_H,\alpha_L)))+p_L(1-\rho\delta_L(\alpha_H,\alpha_L)))^2>0,$$

$$\gamma(\alpha_H,\alpha_L)=-(1-\rho\alpha_L)(1-\Phi(\tau(\alpha_L))),$$

$$\kappa(\alpha_H,\alpha_L)=1+(\alpha_H-\alpha_L)h(\tau(\alpha_L))\tau'(\alpha_L),$$

$$h(\tau(\alpha_L))=\frac{\phi(\tau(\alpha_L))}{1-\Phi(\tau(\alpha_L))}\geq 0.$$

First, from $\gamma(\alpha_H,\alpha_L)$ we note that the waiting cost in the system decreases strictly in α_H for all values of α_L. This implies that at optimum we must have $\alpha_H^*=1$. From $\kappa(\alpha_H,\alpha_L)$ we can see that if $\tau'(\alpha_L)=0$ the waiting cost in the system increases strictly in α_L for all values of α_H. This implies that at optimum we must have $\alpha_L^*=0$. Otherwise, if $\tau'(\alpha_L)<0$, we can see that the FOC for α_L is given by $\kappa(\alpha_L)=1+(1-\alpha_L)h(\tau(\alpha_L))\tau'(\alpha_L)$, and that such FOC increases strictly in α_L: $\kappa'(\alpha_L)=-h(\tau(\alpha_L))\tau'(\alpha_L)+(1-\alpha_L)h'(\tau(\alpha_L))(\tau'(\alpha_L))^2+(1-\alpha_L)h(\tau(\alpha_L))\tau''(\alpha_L)>0$. We can see that $\kappa(1)>0$ and that $\kappa(0)\geq 0$ whenever $1+h(\tau(0))\tau'(0)\geq 0$. Since $\kappa(\alpha_L)$ is strictly increasing in α_L, it follows that whenever $1+h(\tau(0))\tau'(0)\geq 0$, $\alpha_L^*=0$ minimizes the waiting cost in the system. Otherwise, whenever $1+h(\tau(0))\tau'(0)<0$ since $\kappa(\alpha_L)$ is continuous and strictly increasing in α_L, it follows that there is a unique $\alpha_L^*\in(0,1)$ such that $1+(1-\alpha_L^*)h(\tau(\alpha_L^*))\tau'(\alpha_L^*)=0$ that minimizes the waiting cost in the system.

Notice that in this case, we have that all high type customers claim their type and that a proportion equal to $\Phi(\tau(\alpha_L^*))<1$ of low type customers misreport. When the Manager sets the routing prioritization policy $\alpha_H^*=1,\alpha_L^*\in[0,1)$ the prioritization error probabilities are $\delta_H^*=0$ and $\delta_L^*=\alpha_L^*+(1-\alpha_L^*)\Phi(\tau(\alpha_L^*))<1$. It is easy to see from Lemma 1, that this performs better than a FCFS discipline: for a FCFS ($\alpha_H=\alpha_L=1$) we have that $\delta_H=0$ and $\delta_L=1$. Finally, since the routing policies in cases 1 (i.e., $\alpha_H<\alpha_L$) and 2 (i.e., $\alpha_H=\alpha_L$) perform respectively worse/same as a FCFS policy, thus we conclude that the routing policy in this case 3 performs better, and thus arises in equilibrium. ■

Appendix B: Further Details for Experimental Investigation

B.1. Experimental Results per Condition

Condition	$\mathbb{P}$(Claim 5)
1	0.29 ± 0.06
2	0.35 ± 0.06
3	0.34 ± 0.06
4	0.22 ± 0.05
5	0.29 ± 0.06
6	0.29 ± 0.06
7	0.28 ± 0.06
8	0.23 ± 0.05
9	0.24 ± 0.05

Table 2 Proportions and half widths of 95% confidence intervals of participants who reported the die roll 5 across experimental conditions.

B.2. Logistic Regressions

We run logistic regressions where we control for age and gender. We estimate the effects on the probability to claim the number 5 (see Table 3). We conduct likelihood ratio tests between Model (2a) and Model (3a) ($\chi^2 = 0.98$, $df = 1$, p-value $= 0.32$), and Model (2a) and Model (4a) ($\chi^2 = 2.8$, $df = 2$, p-value $= 0.24$) which show that the fit of Model (2a) does not significantly improve by adding the ΔW predictor in Model (3a) and the interaction term in Model (4a). Finally, the likelihood ratio test between Model (2a) and the null model including only an intercept term ($\chi^2 = 13.37$, $df = 3$, p-value 0.004) shows that the fit of Model (2a) is significantly better.

Table 3 Logistic Regressions

	$\mathbb{P}$(Claim 5)			
	(1a)	(2a)	(3a)	(4a)
(Intercept)	-1.03***	-1.19***	-1.29***	-1.09***
	(0.21)	(0.20)	(0.22)	(0.27)
Age	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)
GenderM	0.13	0.13	0.13	0.14
	(0.10)	(0.10)	(0.10)	(0.10)
ΔW	0.01	-	0.01	-0.01
	(0.01)	-	(0.01)	(0.02)
$\Delta\alpha$	-	0.45***	0.45***	0.09
	-	(0.13)	(0.13)	(0.30)
$\Delta\alpha * \Delta W$	-	-	-	0.04
	-	-	-	(0.03)
N	2021	2021	2021	2021
AIC	2398.17	2387.79	2388.81	2388.99
Pseudo R^2	0.00	0.01	0.01	0.01
Pseudo R^2 †	0.01	0.23	0.25	0.25

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

† We compute McFadden's pseudo R^2 at the aggregate experimental condition level. Binary individual responses are grouped at the experimental condition $j \in \{0, \cdots, 9\}$ level with sample size n_j, such that the number of participants that claim the number 5 is captured by a binomial random variable k_j, and the log likelihood is given by $\sum_j log\left(\binom{n_j}{k_j} p_j^{k_j} (1-p_j)^{n_j - k_j}\right)$.

We also run logistic regressions with the covariate α_L instead of $\Delta\alpha$. We are able to do this because $\alpha_H = 1$ was fixed across all experimental conditions such that $\Delta\alpha = 1 - \alpha_L$. The results are presented in Table 4.

Table 4 Logistic Regressions

	$\mathbb{P}$(Claim 5)			
	(1a)	(2a)	(3a)	(4a)
(Intercept)	-1.03***	-0.73***	-0.84***	-1.00***
	(0.21)	(0.19)	(0.22)	(0.25)
Age	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)
GenderM	0.13	0.13	0.13	0.14
	(0.10)	(0.10)	(0.10)	(0.10)
ΔW	0.01	-	0.01	0.03
	(0.01)	-	(0.01)	(0.02)
α_L	-	-0.45***	-0.45***	-0.09
	-	(0.13)	(0.13)	(0.30)
$\alpha_L * \Delta W$	-	-	-	-0.04
	-	-	-	(0.03)
N	2021	2021	2021	2021
AIC	2398.17	2387.79	2388.81	2388.99
Pseudo R^2	0.00	0.01	0.01	0.01
Pseudo R^2 †	0.01	0.23	0.25	0.25

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

† We compute McFadden's pseudo R^2 at the aggregate experimental condition level. Binary individual responses are grouped at the experimental condition $j \in \{0, \cdots, 9\}$ level with sample size n_j, such that the number of participants that claim the number 5 is captured by a binomial random variable k_j, and the log likelihood is given by $\sum_j log\big(\binom{n_j}{k_j} p_j^{k_j} (1-p_j)^{n_j-k_j}\big)$.

Finally, we also run regressions with the covariate $W_L = \alpha_L W_1 + (1-\alpha_L)W_2$, where recall that $W_1 = 2$ min was fixed. In Table 5 we present the results:

Table 5 Logistic Regressions

	$\mathbb{P}$(Claim 5)				
	(1a)	(2a)	(3a)	(4a)	(5a)
(Intercept)	-0.73***	-0.86***	-1.23***	-1.11***	-1.03***
	(0.19)	(0.23)	(0.20)	(0.22)	(0.27)
Age	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
GenderM	0.13	0.13	0.14	0.14	0.14
	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
α_L	-0.45***	-0.45***	-	-	-0.22
	(0.13)	(0.13)	-	-	(0.20)
W_2	-	0.01	-	-0.02	-
	-	(0.01)	-	(0.01)	-
W_L	-	-	0.04***	0.05***	0.03
	-	-	(0.01)	(0.01)	(0.02)
N	2021	2021	2021	2021	2021
AIC	2387.79	2388.81	2386.53	2387.09	2387.32
BIC	2410.24	2416.87	2408.98	2415.15	2415.38
Pseudo R^2	0.00	0.01	0.01	0.01	0.01
Pseudo R^2 †	0.23	0.25	0.22	0.23	0.24

$^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$.

† We compute McFadden's pseudo R^2 at the aggregate experimental condition level. Binary individual responses are grouped at the experimental condition $j\in\{0,\cdots,9\}$ level with sample size n_j, such that the number of participants that claim the number 5 is captured by a binomial random variable k_j, and the log likelihood is given by $\sum_j log\big(\binom{n_j}{k_j}p_i^{k_j}(1-p_j)^{n_j-k_j}\big)$.

First, we notice that the W_L covariate is significant, which is intuitive since it is a function of the upgrading probability, α_L. Also, we notice that in model (4a), where we control for the waiting time in the long queue W_2, the coefficient for such W_2 is negative and not significant. Since $W_L=2\alpha_L+(1-\alpha_L)W_2$, this shows that the majority of explainable lying behaviour can be attributed to the upgrading probability. Now, we observe that the AIC and BIC of model (3a) are lower than that of model (1a), which is consistent with intuition since both models capture the effect of the upgrading probability, but model (3a) in addition captures the effect of the waiting time (even though it is a very small effect). Importantly, we run a Voung test between models (1a) and (3a) for the hypothesis that model (3a) fits the data better than model (1a) ($z = -0.254$, p-value $= 0.4$), which is similar to a Likelihood ratio test but for non-nested models (Vuong 1989) - the 95% confidence intervals for the AIC and BIC difference between models are given by ($-8.462 < AICdiff < 10.982$) and ($-8.462 < BICdiff < 10.982$). This result shows that we cannot reject the null hypothesis that both model fits are equal. Finally, a likelihood ratio test between model (1a) and (5a) ($\chi^2 = 2.47$, $df = 1$, p-value $= 0.12$) shows that the fit of Model (1a) does not significantly improve by adding the W_L predictor. Overall this provides further support for the claim that the great majority of explainable lying behaviour can be attributed to the upgrading probability.

B.3. Robustness Checks

B.3.1. Non-parametric tests: The effect of waiting times. We run a Generalized Cochran-Mantel-Haenszel test (Agresti 2003) for the conditional association between the proportion of participants that reported the number 5 and the ΔW condition, conditional on the levels of the $\Delta\alpha$ condition. We find that the conditional association is not significant ($M^2 = 1.36$, $df = 2$, p-value = 0.51). We also run a Chi-square test for the marginal association between the proportion of participants that reported the number 5 and the ΔW condition. We find that the marginal association is not significant ($\chi^2 = 1.22$, $df = 2$, p-value = 0.54). Moreover, for each $\Delta\alpha$ level (i.e., 0.1, 0.5, 1), we run Jonckheere–Terpstra tests for an order, both ascending and descending. We find that the proportion of participants that claim the number 5 does not significantly increase monotonically (p-values = $0.78, 0.11, 0.20$), nor decrease monotonically (p-values = $0.22, 0.89, 0.79$) as ΔW increases in each of the respective $\Delta\alpha$ levels. Also, by aggregating the data (ignoring the $\Delta\alpha$ levels), both the ascending (p-value $= 0.24$) and descending (p-value 0.75) trends are not significant in the Jonckheere–Terpstra test.

B.3.2. Non-parametric tests: The effect of routing probabilities. We run a Generalized Cochran-Mantel-Haenszel test for the conditional association between the proportion of participants that reported the number 5 and the $\Delta\alpha$ condition, conditional on the levels of ΔW. We find that the conditional association is significant ($M^2 = 12.16$, $df = 2$, p-value = 0.002). We also run a Chi-square test for the marginal association between the proportion of participants that reported the number 5 and the $\Delta\alpha$ condition. We find that the marginal association is significant ($\chi^2 = 12.03$, $df = 2$, p-value $= 0.002$). For each ΔW level (i.e., 3 min, 8 min, 13 min), we run Jonckheere–Terpstra tests for an order, both ascending and descending. We find that the proportion of participants who claim the number 5 does not significantly decrease monotonically (p-values $= 0.61, 0.99, 0.97$) as the $\Delta\alpha$ increases in any of the respective ΔW levels. Moreover, while that proportion does not significantly increase monotonically for the $2-5$ min level (p-value $= 0.39$), it does significantly decrease monotonically in the higher wait-time conditions, i.e., $2-10$ min (p-value $= 0.009$) and $2-15$ min (p-value $= 0.03$). Also, by aggregating the data (ignoring the ΔW levels), the proportion of participants that claim 5 significantly increases monotonically as the $\Delta\alpha$ increases (p-value 0.005).

B.3.3. Simulations. In our logistic regressions, we study the reporting behaviour of participants (i.e., whether they report a number 5) in order to derive conclusions about their misreporting behaviour. We found a significant $\Delta\alpha$ effect on the reporting behaviour. We concluded that this is attributed to changes in the lying behaviour. Since we do not observe individual die outcomes in the experiment, the noise in the realizations of the die rolls may lead to a wrong conclusion that $\Delta\alpha$ has an effect on lying behaviour, when in fact there is no effect, i.e., there may be a type 1 error with respect to the $\Delta\alpha$ condition.

To understand the impact of the sampling variation of the die outcome on the type 1 error probability of the $\Delta\alpha$ condition, we conduct a simulation analysis, where for each of the $j = 9$ experimental conditions (with sample size N_j), we simulate N_j independent Bernoulli(1/6) rolls. This captures the fact that participants get a number 5 with probability 1/6, and any other number with probability 5/6. Then, based on the actual die outcome, we simulate participants reports in the following way: If the actual die outcome is 5 we simulate a report equal to 5 with probability 1, and if the actual die outcome is not 5 we simulate a report equal to 5 with probability $\mathbb{P}(misreport)$, where we select different values of $\mathbb{P}(misreport)$ as a further robustness check. Importantly, to investigate the type 1 error, we assume that the lying probability in our simulations does not depend on the experimental conditions. We consider $\mathbb{P}(misreport) = 0$, $\mathbb{P}(misreport) = 0.14$ (which is the average lying rate in our data set [9]), and $\mathbb{P}(misreport) = 0.35$ (which is the maximum possible lying rate in our data set, and assumes that no participant actually rolled a number 5). Finally, once a data set is generated

[9] We use the method proposed in Hugh-Jones (2019) to estimate lying rates in binary lying games. The technique estimates the excess number of 5s reports to what would be expected under full honesty based on a bayesian approach.

based on the aforementioned process, we fit a logistic regression with linear predictor $\beta_0 + \beta_1 \Delta W + \beta_2 \Delta\alpha$. For a total of 10,000 simulation runs, we use the p-values associated with the coefficient β_2 to compute the proportion of runs where such a coefficient is significant at different levels (i.e., $0.01, 0.05, 0.01, 0.001$), that is, where there is a type 1 error. The following table shows the results of our simulation analysis.

Table 6 Simulation Results: Type 1 error probability

Significance level	$\mathbb{P}(misreport)$ 0	0.14	0.35
0.1	0.1019	0.0973	0.0969
0.05	0.0501	0.0508	0.0508
0.01	0.0097	0.0094	0.0087
0.001	0.0006	0.0007	0.0017

Overall, based on our simulation results we corroborate that the probability of a type 1 error is very close to the considered significance levels. Thus, we can confidently conclude that the observed differences in the reporting behaviour can be attributed to differences in the lying behaviour, and not to sampling variation. Finally, we want to mention that, in a related exercise, Fries et al. (2021) run a similar analysis in which they compare the frequency of reported 5s with the expected probability of success of 16.66% using binomial tests. They also corroborate that the type 1 error is close to the considered significance levels. Similarly, Fries and Parra (2021) show that frequentist methods in the die roll game produce accurate confidence intervals for the average die roll when the sample is about 100 observations or more.

B.4. Structural Estimation

We estimate the relevant lying aversion parameters when considering different distributions. In particular, we investigate an exponential distribution with density $\phi(\theta)=\gamma e^{-\gamma\theta}$, a uniform distribution with density $\phi(\theta)=1/\bar{\theta}$, a half-normal distribution with density $\phi(\theta)=\frac{2\sigma}{\pi}e^{-\frac{(\theta\sigma)^2}{\pi}}$, and a half logistic distribution with density $\phi(\theta)=\frac{2e^{\theta/\sigma_l}}{\sigma_l(1+e^{\theta/\sigma_l})^2}$. We do the estimation considering different sensible lying cost specifications to investigate how well they fit the experimental data. In particular, we are interested to show that even a very simple specification for $\tau(\boldsymbol{\alpha})$ as per our proposed model in §6.4 can achieve better fit in comparison to the sensible lying costs specifications from the literature. For this, we propose the specification $\tau(\boldsymbol{\alpha})=2-\alpha_L$, which indeed complies with the assumptions developed in §6.4.1.

Let $\mathbf{X}=\{y^i_j,\alpha^j_L,W^j_2\}$ represent our experimental data set, where y^i_j denotes the observed indicator random variate for participant i ($y^i_j=1$ if participant i reports 5), in an experimental condition j with upgrading probability α^j_L and waiting time in the long queue W^j_2. We assume that Y^i_j are mutually independent across participants. We define $p_j:=\mathbb{P}(Y^i_j=1)$ which depends on the actual die roll for participants as follows: $\mathbb{P}(Y^i_j=1|\text{actual die outcome}=5)=1$ and $\mathbb{P}(Y^i_j=1|\text{actual die outcome}\neq 5)=\Phi(f(\boldsymbol{\alpha},\mathbf{W}))$ where Φ is the cumulative distribution function of the lying aversion, Θ, and $f(\boldsymbol{\alpha},\mathbf{W})$ is some function that depends on the considered lying cost specification (see Table 7 below).

Let β represent the relevant lying aversion parameter to estimate depending on the considered distribution. Given the experimental design, and under the assumption that participants roll a fair die, we obtain that: $p_j(\beta)=\mathbb{P}(Y^i_j=1)=\frac{1}{6}\mathbb{P}(Y^i_j=1|\text{actual die outcome}=5)+\frac{5}{6}\mathbb{P}(Y^i_j=1|\text{actual die outcome}\neq 5)=\frac{1}{6}+\frac{5}{6}\Phi(f(\boldsymbol{\alpha},\mathbf{W}))$. Also, for our estimation we group binary individual responses at the experimental condition $j\in\{0,\cdots,9\}$ level with sample size n_j, such that the number of participants who claim the number 5 is captured by a binomial random variable k_j, and the log-likelihood is given by

$$\mathcal{L}(\beta|\mathbf{X})=\sum_j log\Big(\binom{n_j}{k_j}p_j^{k_j}(\beta)(1-p_j(\beta))^{n_j-k_j}\Big). \tag{14}$$

We use MLE to estimate the relevant lying aversion parameter that maximizes $\mathcal{L}$ for different sensible lying cost specifications. Table 7 summarizes the estimation results and the model fit for all considered lying cost specifications. We can see, based on the Akaike information criterion (AIC), that the proposed specification for $\tau(\boldsymbol{\alpha})$ as per our proposed model in §6.4 can achieve a better fit in comparison to the sensible lying costs specifications from the literature.

Table 7 Structural estimation results for different sensible lying cost specifications.

Lying Cost	Predicted Misreporting	$\Theta \sim Exp(\gamma)$		$\Theta \sim Unif(0, \bar{\theta})$		$\Theta \sim HalfNormal(\sigma)$		$\Theta \sim HalfLogis(\sigma_l)$	
Specification	Probability	γ	AIC	$\bar{\theta}$	AIC	σ	AIC	σ_l	AIC
Fixed	$\Phi((\alpha_H - \alpha_L)(W_2 - W_1))$	0.03*** (0.003)	84.9	43.05*** (4.063)	89	33.48*** (3.267)	88.25	20.83*** (2.05)	88
Linear	$\Phi(1)$	0.15*** (0.014)	72	7.41*** (0.658)	72	5.88*** (0.527)	72	3.68*** (0.331)	72
Quadratic	$\Phi(((\alpha_H - \alpha_L)(W_2 - W_1))^{-1})$	0.07*** (0.013)	180.5	15.73*** (2.515)	183	12.31*** (2.009)	182.7	7.66*** (1.257)	182.7
Proposed	$\Phi(2 - \alpha_L)$	0.01*** (0.009)	61.4	11.21*** (0.968)	61	8.9*** (0.778)	61.1	5.57*** (0.488)	61.11

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

† Values shown are the MLE with associated (standard error).

Appendix C: Additional Experimental Investigation

For this investigation, we use the same experimental procedure, and criteria for the recruitment of participants and exclusions, as in Section §6, so we do not repeat these here. Based on the results in our main experimental investigation, in here we are interested in understanding the relative effect of α_H and α_L on the misreporting probability (8) to gain further insight into $\tau(\boldsymbol{\alpha})$. Formally, we set the following null hypotheses:

H1.b. Changing the routing probability α_H does not have an effect on the proportion of participants that misreport.

H2.b. Changing the routing probability α_L does not have an effect on the proportion of participants that misreport.

Based on the above, in order to test these hypotheses, participants were randomly assigned to one of nine experimental conditions (see Table 8) which differ in both the routing probabilities, α_H and α_L, while keeping the waiting times in the queues constant.

Condition	α_H	α_L	W_1	W_2	Sample
1	1	0	2 min	15 min	141
2	1	0.2	2 min	15 min	140
3	1	0.4	2 min	15 min	141
4	0.8	0	2 min	15 min	138
5	0.8	0.2	2 min	15 min	141
6	0.8	0.4	2 min	15 min	133
7	0.6	0	2 min	15 min	138
8	0.6	0.2	2 min	15 min	140
9	0.6	0.4	2 min	15 min	130

Table 8 Experimental conditions.

We set the target sample size for the experiment and our analysis plans a priori. We pre-registered our experiment, and the corresponding As Predicted document can be found at: `https://aspredicted.org/2DV_J7D`. A total of 1,633 participants (41.58% female, mean age $M_{age} = 34.25$, standard deviation $SD_{age} = 11.70$) were recruited. In our experiment, we have a completion rate of 87%. From those who completed the experiment, 87% of participants took, on average, less than 30 seconds to click the advance in queue buttons, and the mean and median average click time was 17 seconds and 11 seconds, respectively. After exclusions, we are left with a sample of 1,242 participants (43.96% female, mean age $M_{age} = 34.78$, standard deviation $SD_{age} = 10.89$).

C.1. Experimental Results

In Figure 2, we present the proportion of participants who reported the number 5 across experimental conditions. In Table 9, we present corresponding proportions. Participants have the incentive to report the number 5 to reduce their waiting time in the queue. We run exact binomial tests for the proportions of participants who reported the number 5 in each condition compared to the proportion that would have reported the number 5 under full honesty i.e., 1/6, (p-values for conditions 1 to 9 = (0.49, 0.01, 0.24, 0.07, 0.06, 0.15, 0.01, 0.12, 0.32)). We can see that, in some conditions, participants might have been fully honest. We note that these results provide further support for the existence of lying costs.

Condition	$\mathbb{P}$(Claim 5)
1	0.17 ± 0.06
2	0.24 ± 0.07
3	0.19 ± 0.06
4	0.22 ± 0.07
5	0.22 ± 0.07
6	0.20 ± 0.07
7	0.24 ± 0.07
8	0.21 ± 0.07
9	0.18 ± 0.07

Table 9 Proportions and half widths of 95% confidence intervals of participants who reported the die roll 5 across experimental conditions.

We run logistic regressions where we control for age and gender. We estimate the treatments effects on the probability to claim the number 5 (see Table 10). First, in terms of α_H, we can see that in all model specifications the coefficients of the α_H covariate are negative. This is consistent with the hypothesised effect

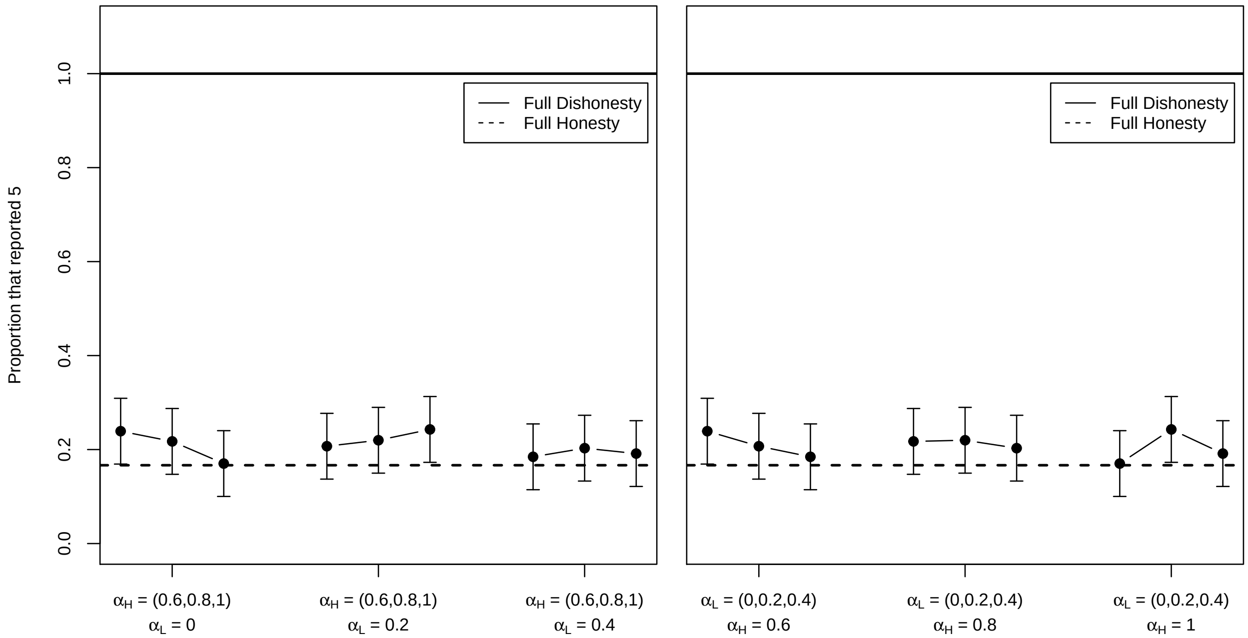

Figure 2 Proportions of participants that reported the number 5 across experimental conditions.

in (Celse et al. 2019), however, such coefficients are not significant. We are not able to reject H1.a. Similarly, in terms of α_L, we can see that in all model specifications the coefficients of the α_L covariate are negative. This is consistent with the results in our main experimental investigation, however, such coefficients are not significant. In light of the significant results in our main experimental investigation, we believe that the lack of significance in this study arises from the fact that overall, participants were too honest in this experiment, such that effects were not significantly identified. We want to mention that in Proposition 1, our theoretical results show that for a population as the one observed in this experiment, it is optimal to implement an honour system.

Table 10 Logistic Regressions

	$\mathbb{P}$(Claim 5)			
	(1a)	(2a)	(3a)	(4a)
(Intercept)	-1.25 **	-1.31 ***	-1.20 **	-0.77
	(0.42)	(0.26)	(0.43)	(0.58)
Age	0.00	0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.01)
GenderM	0.00	0.00	0.00	0.01
	(0.14)	(0.14)	(0.14)	(0.14)
α_H	-0.14	-	-0.14	-0.70
	(0.43)	-	(0.43)	(0.67)
α_L	-	-0.23	-0.23	-2.53
	-	(0.43)	(0.43)	(2.15)
$\alpha_H * \alpha_L$	-	-	-	2.86
	-	-	-	(2.63)
N	1242	1242	1242	1242
AIC	1279.70	1279.52	1281.41	1282.23
Pseudo R^2	0.00	0.00	0.00	0.00
Pseudo R^2 †	0.01	0.01	0.01	0.05

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

† We compute McFadden's pseudo R^2 at the aggregate experimental condition level. Binary individual responses are grouped at the experimental condition $j \in \{0, \cdots, 9\}$ level with sample size n_j, such that the number of participants that claim the number 5 is captured by a binomial random variable k_j, and the log likelihood is given by $\sum_j log\big(\binom{n_j}{k_j} p_j^{k_j}(1-p_j)^{n_j-k_j}\big)$.

Appendix D: General Conditions for the Optimal Scheduling Policy

We consider the case in which $\alpha_H > \alpha_L$, such that L type customers are those that have the incentive to misreport. Based on this, to derive general conditions for the identified optimal scheduling policy, we assume that H type customers never down-report, which is intuitive, and that L type customers best respond by misreporting with some probability that does not depend on waiting times. This is consistent with our experimental results and with the extant literature that finds that misreporting behaviour is insensitive to material incentives (Abeler et al. 2019). The insensitivity to waiting times implies that customers' misreporting equilibrium is a function that depends only on the routing probabilities, $\eta(\alpha_H, \alpha_L)$. To see this, note that, generally speaking, the misreporting equilibrium probability (call it ν) is given by the solution of the fixed-point problem $\nu = \eta(\alpha_H, \alpha_L, W_1(\nu), W_2(\nu))$ (where waiting times and misreporting are consistent and endogenously determined and η denotes the best-response function). However, since we have that $\eta(\alpha_H, \alpha_L)$, it simply follows that $\nu = \eta(\alpha_H, \alpha_L)$. This is intuitive since customers individual best responses do not care for the waiting times. Since customers affect each other only through the resulting waiting times, all individual best responses represent a best response correspondence, forming a customer equilibrium where a proportion equal to $\eta(\alpha_H, \alpha_L)$ misreports. Based on this, the task at hand is to determine under which assumptions regarding the structure of $\eta(\alpha_H, \alpha_L)$, our identified optimal scheduling policy in the main paper will hold. Here, we present a general set of sufficient conditions.

PROPOSITION 2. *Assume that $\alpha_H > \alpha_L$, that H type customers never down-report, that L type customers misreport with probability $\eta(\alpha_H, \alpha_L)$, and that $\eta(\alpha_H, \alpha_L)$ is differentiable at all $\alpha_H > \alpha_L$. If $-\frac{\partial}{\partial \alpha_L}\eta(\alpha_H, \alpha_L) \geq \frac{\partial}{\partial \alpha_H}\eta(\alpha_H, \alpha_L)$ for all $\alpha_H \geq \alpha_L$, then it follows that the optimal scheduling policy is given by $(\alpha_H^* = 1, \alpha_L^* = 0)$ whenever $-\frac{\partial}{\partial \alpha_L}\eta(1,0) \leq 1 - \eta(1,0)$, and $(\alpha_H^* = 1, \alpha_L^* \in (0,1))$ otherwise.*

Proposition 2, shows that if the marginal effect of upgrading to mitigate misreporting is greater than or equal to the marginal effect of downgrading (for all $\alpha_H > \alpha_L$), then the same structure for the optimal scheduling policy (i.e., honour system and upgrading policy) holds. Importantly, Proposition 2 allows to directly test if a particular customer problem will give rise to upgrading. For example, consider the customer problem (2). In this case, since $\alpha_H > \alpha_L$, we know that H type customers do not down-report. Now, assume that L type customers experience a lying cost captured by $c_L(W_2 - W_1)$, that is by the difference between expected waiting times in *priority queues* (rather than in *priority classes* as in the main paper).

In this case, we have that $\eta = \Phi(\alpha_H - \alpha_L)$. We can see that the conditions of the above proposition are satisfied: $-\frac{\partial}{\partial \alpha_L}\eta(\alpha_H,\alpha_L) = \phi(\alpha_H - \alpha_L) \geq \phi(\alpha_H - \alpha_L) = \frac{\partial}{\partial \alpha_H}\eta(\alpha_H,\alpha_L)$. Based on this, Proposition 2 would prescribe an honour system whenever $\phi(1) \leq 1 - \Phi(1) \iff h(1) \leq 1$, and an upgrading policy otherwise. For another example, assume that L type customers experience a lying cost captured by $c_L(W_2 - W_H)$, that is by the difference between the *deserved* waiting cost and the expected one from misreporting. In this case we have that $\eta = \Phi(\frac{\alpha_H - \alpha_L}{\alpha_H})$. We can see that the conditions of the above proposition are satisfied: $-\frac{\partial}{\partial \alpha_L}\eta(\alpha_H,\alpha_L) = \phi(\frac{\alpha_H-\alpha_L}{\alpha_H})\frac{1}{\alpha_H} \geq \phi(\frac{\alpha_H-\alpha_L}{\alpha_H})\frac{\alpha_L}{\alpha_H^2} = \frac{\partial}{\partial \alpha_H}\eta(\alpha_H,\alpha_L)$. Based on this, Proposition 2 would prescribe an honour system whenever $\phi(1) \leq 1 - \Phi(1) \iff h(1) \leq 1$, and an upgrading policy otherwise. Finally, consider a different formalization of the customer problem, based on a discrete choice framework. In this case, the misreporting choice probability can be modelled by $\mathbb{P}(\varepsilon < \beta_0 + \beta_1\Delta\alpha + \beta_2\Delta W)$, where ε is a random term, and $\beta_0 + \beta_1\Delta\alpha + \beta_2\Delta W$ is a linear predictor that captures differences across choice alternatives (Train 2009). If we let ε be logistically distributed, we obtain a logit choice model. Notice that our logistic regressions in Table 3 suggest that the misreporting choice probability can be well captured by $\mathbb{P}(\varepsilon < \beta_0 + \beta_1\Delta\alpha)$, such that the misreporting probability is equal to $\Psi(\beta_0 + \beta_1\Delta\alpha)$, where Ψ is the CDF of the logistic distribution. Based on this, we can see that the conditions of the above proposition are satisfied: $-\frac{\partial}{\partial \alpha_L}\eta(\alpha_H,\alpha_L) = \psi(\beta_0 + \beta_1\Delta\alpha)\beta_1 \geq \psi(\beta_0 + \beta_1\Delta\alpha)\beta_1 = \frac{\partial}{\partial \alpha_H}\eta(\alpha_H,\alpha_L)$, where ψ is the logistic density function. Based on this, and assuming that H types never down-report, Proposition 2 would prescribe an honour system whenever $\psi(\beta_0 + \beta_1)\beta_1 \leq 1 - \Psi(\beta_0 + \beta_1) \iff h(\beta_0 + \beta_1)\beta_1 \leq 1$, and an upgrading policy otherwise. We now present the proof of the above Proposition 2.

Proof. In this proof, for any function $q(\cdot)$, we use $q'_z(\cdot)$ to denote the first partial derivative of $q(\cdot)$ with respect to z. Consider the case in which customers misreport with probability $\eta(\alpha_H,\alpha_L)$, where $\alpha_H \geq \alpha_L$. Based on the above, the Manager anticipates the following prioritization error probabilities:

$$\delta_H(\alpha_H,\alpha_L) = 1 - \alpha_H,$$

$$\delta_L(\alpha_H,\alpha_L) = \alpha_L + (\alpha_H - \alpha_L)\eta(\alpha_H,\alpha_L).$$

The Manager defines the routing probabilities $\alpha_H, \alpha_L \in [0,1]$ in order to minimize the expected waiting cost C_s in the system which can be written as:

$$C_s = c \cdot f(\alpha_H,\alpha_L),$$

where

$$c = \left(\frac{\rho^2}{1-\rho}\right) > 0,$$

$$f(\alpha_H,\alpha_L) = \frac{p_H c_H(1-\rho(1-\delta_H(\alpha_H,\alpha_L))) + p_L c_L(1-\rho\delta_L(\alpha_H,\alpha_L))}{p_H(1-\rho(1-\delta_H(\alpha_H,\alpha_L))) + p_L(1-\rho\delta_L(\alpha_H,\alpha_L))},$$

$$f'_{\alpha_H} = \frac{d(\alpha_H,\alpha_L)\gamma(\alpha_H,\alpha_L)}{g(\alpha_H,\alpha_L)},$$

$$f'_{\alpha_L} = \frac{d(\alpha_H,\alpha_L)(1-\rho\alpha_H)\kappa(\alpha_H,\alpha_L)}{g(\alpha_H,\alpha_L)},$$

$$d(\alpha_H,\alpha_L) = p_H p_L \rho(c_H - c_L) > 0,$$

$$g(\alpha_H,\alpha_L) = (p_H(1-\rho(1-\delta_H(\alpha_H,\alpha_L))) + p_L(1-\rho\delta_L(\alpha_H,\alpha_L)))^2 > 0,$$

$$\gamma(\alpha_H,\alpha_L) = -(1-\rho\alpha_L)(1-\eta(\alpha_H,\alpha_L)) + (1-\rho\alpha_H)(\alpha_H-\alpha_L)\eta'_{\alpha_H}(\alpha_H,\alpha_L),$$

$$\kappa(\alpha_H,\alpha_L) = (1-\eta(\alpha_H,\alpha_L)) + (\alpha_H-\alpha_L)\eta'_{\alpha_L}(\alpha_H,\alpha_L),$$

We want to solve the following constraint minimization problem:

$$Min \;\; c\cdot f(\alpha_H,\alpha_L)$$

$$s.t$$

$$\alpha_L \geq 0 \iff -\alpha_L \leq 0$$

$$\alpha_L \leq \alpha_H \iff \alpha_L - \alpha_H \leq 0$$

$$\alpha_H \leq 1 \iff \alpha_H - 1 \leq 0$$

We construct the lagrangian:

$$\mathcal{L} = cf(\alpha_H,\alpha_L) - \mu_1(\alpha_L) + \mu_2(\alpha_L-\alpha_H) + \mu_3(\alpha_H - 1),$$

and derive the KKT conditions:

Stationarity

$$\mathcal{L}'_{\alpha_H} = cf'_{\alpha^*_H} - \mu_2 + \mu_3 = 0,$$

$$\mathcal{L}'_{\alpha_L} = cf'_{\alpha^*_L} - \mu_1 + \mu_2 = 0.$$

Complementary Slackness

$$\mu_1(-\alpha_L^*) = 0,$$

$$\mu_2(\alpha_L^* - \alpha_H^*) = 0,$$

$$\mu_3(\alpha_H^* - 1) = 0.$$

Dual Feasibility

$$\mu_1, \mu_2, \mu_3 \geq 0.$$

Primal Feasibility

$$\alpha_L^* \geq 0,$$

$$\alpha_L^* \leq \alpha_H^*,$$

$$\alpha_H^* \leq 1.$$

Potential Candidate 1: $(\alpha_H^* = 1, \alpha_L^* = 0)$.

Complementary Slackness

$\mu_1 \neq 0 \Rightarrow \alpha_L^* = 0,$

$\mu_2 = 0,$

$\mu_3 \neq 0 \Rightarrow \alpha_H^* = 1.$

Stationarity

$\mu_3 = -cf'_{\alpha_H^*} = -\frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)} = -\frac{cd(1,0)\gamma(1,0)}{g(1,0)},$
$\mu_1 = cf'_{\alpha_L^*} = \frac{cd(\alpha_H^*,\alpha_L^*)(1-\rho\alpha_H^*)\kappa(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)} = \frac{cd(1,0)(1-\rho)\kappa(1,0)}{g(1,0)}.$

Dual Feasibility

$\mu_3 \geq 0 \iff \frac{cd(1,0)\gamma(1,0)}{g(1,0)} \leq 0 \iff \gamma(1,0) \leq 0 \iff \eta'_{\alpha_H}(1,0) \leq (1-\eta(1,0))/(1-\rho),$
$\mu_1 \geq 0 \iff \frac{cd(1,0)(1-\rho)\kappa(1,0)}{g(1,0)} \geq 0 \iff \kappa(1,0) \geq 0 \iff -\eta'_{\alpha_L}(1,0) \leq (1-\eta(1,0)).$

Primal Feasibility

$\alpha_L^* \geq 0,$

$\alpha_L^* \leq \alpha_H^*,$

$\alpha_H^* \leq 1.$

Notice that as long as $-\eta'_{\alpha_L}(1,0) \geq (1-\rho)\eta'_{\alpha_H}(1,0)$, this solution is a possible candidate whenever $-\eta'_{\alpha_L}(1,0) \leq (1-\eta(1,0))$.

Potential Candidate 2: $(\alpha_H^* = 1, \alpha_L^* \in (0,1))$.

Complementary Slackness

$\mu_1 = 0,$

$\mu_2 = 0,$

$\mu_3 \neq 0 \Rightarrow \alpha_H^* = 1.$

Stationarity

$cf'_{\alpha_L^*} = 0 \iff \frac{cd(\alpha_H^*,\alpha_L^*)(1-\rho\alpha_H^*)\kappa(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)} = \frac{cd(1,\alpha_L^*)(1-\rho)\kappa(1,\alpha_L^*)}{g(1,\alpha_L^*)} = 0 \iff \kappa(1,\alpha_L^*) = 0 \iff -\eta'_{\alpha_L}(1,\alpha_L^*)) = \frac{(1-\eta(1,\alpha_L^*))}{1-\alpha_L^*},$

$\mu_3 = -cf'_{\alpha_H^*} = -\frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)} = -\frac{cd(1,\alpha_L^*)\gamma(1,\alpha_L^*)}{g(1,\alpha_L^*)}.$

Dual Feasibility

$\mu_3 \geq 0 \iff \frac{cd(1,\alpha_L^*)\gamma(1,\alpha_L^*)}{g(1,\alpha_L^*)} \leq 0 \iff \gamma(1,\alpha_L^*) \leq 0 \iff -\eta'_{\alpha_L}(1,\alpha_L^*) \geq \frac{(1-\rho)\eta'_{\alpha_H}(1,\alpha_L^*)}{1-\rho\alpha_L^*}.$

Notice that $\frac{1-\rho}{1-\rho\alpha_L^*} < 1$.

Primal Feasibility

We define the function $v(\alpha_L) = -\eta'_{\alpha_L}(1,\alpha_L)(1-\alpha_L) - (1-\eta(1,\alpha_L))$. We can see that $v(0) > 0$ whenever $-\eta'_{\alpha_L}(1,0) > 1-\eta(1,0)$, and that $v(1) < 0$. Since $v(\alpha_L)$ is continuous in $\alpha_L \in (0,1)$ it follows that there is at least one $\alpha_L^* \in (0,1)$ such that $v(\alpha_L^*) = 0$. To preserve generality in our results, we do not impose any further restrictions to *guarantee* the uniqueness of α_L^*. The purpose of the above is to show that whenever $-\eta'_{\alpha_L}(1,0) > 1-\eta(1,0)$, upgrading is optimal is a candidate.

$\alpha_L^* \geq 0 \iff -\eta'_{\alpha_L}(1,0) > 1 - \eta(1,0)$

$\alpha_L^* \leq \alpha_H^*,$

$\alpha_H^* \leq 1.$

Potential Candidate 3: $(\alpha_H^* \in (\alpha_L^*, 1), \alpha_L^* = 0)$.

Complementary Slackness

$\mu_1 \neq 0 \Rightarrow \alpha_L^* = 0,$

$\mu_2 = 0,$

$\mu_3 = 0,$

Stationarity

$cf'_{\alpha_H^*} = 0 \iff \frac{cd(\alpha_H^*, \alpha_L^*)\gamma(\alpha_H^*, \alpha_L^*)}{g(\alpha_H^*, \alpha_L^*)} = 0 \iff \frac{cd(\alpha_H^*, 0)\gamma(\alpha_H^*, 0)}{g(\alpha_H^*, 0)} = 0 \iff \gamma(\alpha_H^*, 0) = 0 \iff \eta'_{\alpha_H}(\alpha_H^*, 0) = \frac{1-\eta(\alpha_H^*, 0)}{(1-\rho\alpha_H^*)\alpha_H^*},$

$\mu_1 = cf'_{\alpha_L^*} = \frac{cd(\alpha_H^*, \alpha_L^*)(1-\rho\alpha_H^*)\kappa(\alpha_H^*, \alpha_L^*)}{g(\alpha_H^*, \alpha_L^*)} = \frac{cd(\alpha_H^*, 0)(1-\rho\alpha_H^*)\kappa(\alpha_H^*, 0)}{g(\alpha_H^*, 0)}.$

Dual Feasibility

$\mu_1 \geq 0 \iff \frac{cd(\alpha_H^*, 0)(1-\rho\alpha_H^*)\kappa(\alpha_H^*, 0)}{g(\alpha_H^*, 0)} \geq 0 \iff \kappa(\alpha_H^*, 0) \geq 0 \iff -\eta'_{\alpha_L}(\alpha_H^*, 0) \leq (1 - \rho\alpha_H^*)\eta'_{\alpha_H}(\alpha_H^*, 0).$

Notice that as long as $-\eta'_{\alpha_L}(\alpha_H^*, 0) \geq \eta'_{\alpha_H}(\alpha_H^*, 0)$, this does not hold.

Potential Candidate 4: $(\alpha_H^* \in (\alpha_L^*, 1), \alpha_L^* \in (0, \alpha_H^*))$.

Complementary Slackness

$\mu_1 = 0,$

$\mu_2 = 0,$

$\mu_3 = 0.$

Stationarity

$cf'_{\alpha_L^*}=0 \iff \frac{cd(\alpha_H^*,\alpha_L^*)(1-\rho\alpha_H^*)\kappa(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)}=0 \iff \kappa(\alpha_H^*,\alpha_L^*)=0 \iff 1-\eta(\alpha_H^*,\alpha_L^*)=-\eta'_{\alpha_L}(\alpha_H^*,\alpha_L^*)(\alpha_H^*-\alpha_L^*)$,

$cf'_{\alpha_H^*}=0 \iff \frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)}=0 \iff \frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)}=0 \iff \gamma(\alpha_H^*,\alpha_L^*)=0 \iff 1-\eta(\alpha_H^*,\alpha_L^*)=\eta'_{\alpha_H}(\alpha_H^*,\alpha_L^*)(\alpha_H^*-\alpha_L^*)\frac{(1-\rho\alpha_L^*)}{(1-\rho\alpha_H^*)}$.

Since we have that $\alpha_H^*>\alpha_L^*$, we can see that this can never happen if $-\eta'_{\alpha_L}(\alpha_H^*,\alpha_L^*)\geq\eta'_{\alpha_H}(\alpha_H^*,\alpha_L^*)$.

Potential Candidate 5: $(\alpha_H^*=\alpha_L^*\in(0,1))$.

Complementary Slackness

$\mu_1=0$,

$\mu_2\neq 0\Rightarrow \alpha_H^*=\alpha_L^*$,

$\mu_3=0$.

Stationarity

$\mu_2=cf'_{\alpha_H^*}=\frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)}=-\frac{cd(\alpha_H^*,\alpha_L^*)(1-\rho\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)}$.

Dual Feasibility

$\mu_2\geq 0 \iff -\frac{cd(\alpha_H^*,\alpha_L^*)(1-\rho\alpha_H^*)}{g(\alpha_H^*,\alpha_L^*)}\geq 0 \iff -(1-\rho\alpha_L^*)\geq 0$ which is never the case.

Potential Candidate 6: $(\alpha_H^*=\alpha_L^*=0)$.

Complementary Slackness

$\mu_1\neq 0\Rightarrow \alpha_L^*=0$,

$\mu_2\neq 0\Rightarrow \alpha_H^*=\alpha_L^*$,

$\mu_3=0$.

Stationarity

$\mu_2=cf'_{\alpha_H^*}=\frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)}=-\frac{cd(0,0)}{g(0,0)}$.

Dual Feasibility

$\mu_2 \geq 0 \iff -\frac{cd(0,0)}{g(0,0)} \geq 0$ which is never the case.

Potential Candidate 7: $(\alpha_H^* = \alpha_L^* = 1)$.

Complementary Slackness

$\mu_1 = 0,$

$\mu_2 \neq 0 \Rightarrow \alpha_H^* = \alpha_L^*,$

$\mu_3 \neq 0 \Rightarrow \alpha_H^* = 1.$

Stationarity

$\mu_2 = -cf'_{\alpha_L^*} = -\frac{cd(\alpha_H^*,\alpha_L^*)\gamma(\alpha_H^*,\alpha_L^*)}{g(\alpha_H^*,\alpha_L^*)} = -\frac{cd(1,1)(1-\rho)}{g(1,1)}.$

Dual Feasibility

$\mu_2 \geq 0 \iff -\frac{cd(1,1)(1-\rho)}{g(1,1)} = -(1-\rho) \geq 0$ which is never the case.

We can see that if $-\eta'_{\alpha_L}(\alpha_H,\alpha_L) \geq \eta'_{\alpha_H}(\alpha_H,\alpha_L)$, only candidates 1 and 2 comply with the KKT conditions. In particular, we notice that candidate 1 (i.e., $\alpha_H^* = 1, \alpha_L^* = 0$) holds whenever $-\eta'_{\alpha_L}(1,0) \leq 1 - \eta(1,0)$ whereas candidate 2 (i.e., $\alpha_H^* = 1, \alpha_L^* \in (0,1)$ holds whenever $-\eta'_{\alpha_L}(1,0) > 1 - \eta(1,0)$. Since these two candidates are mutually exclusive, we conclude that the KKT necessary conditions are also sufficient. ■