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

Persuasive Communication in Social Service Operations

Arturo Estrada Rodriguez, Rouba Ibrahim
School of Management, University College London, 1 Canada Square, London E14 5AB
arturo.rodriguez.18@ucl.ac.uk , rouba.ibrahim@ucl.ac.uk

Mirko Kremer
Management Department, Frankfurt School of Finance and Management, 60314 Frankfurt am Main, Germany
m.kremer@fs.de

We study the effectiveness of information design as a managerial lever to mitigate the overuse of critical resources in congestion-prone social service systems. Leveraging the service provider's informational advantage about relevant aspects of the system, effective communication requires the sharing of carefully curated information to persuade low-need customers to forego service for the benefit of customers with higher service needs. To study whether effective communication can arise in equilibrium, we design controlled laboratory experiments to test the predictions of a queueing-game theoretic model that endogenizes the implementation of information-sharing policies. Our main result is that communication increases social welfare even when the service provider lacks the ability to formally commit to their information policy (as usually is the case in practical settings), i.e., under conditions where standard theory predicts that communication fails because it lacks credibility and thus fails to affect customer behaviour.

Key words: communication; persuasion; queueing; commitment; noisy decision-making; lying aversion

1. Introduction

"We are experiencing long waits"; "This item is almost out of stock"; "Your driver is almost there";

Service providers communicate to their customers real-time information about aspects of their operational systems and processes that are otherwise unobservable for the customers. Such communication, despite its occasional vagueness,[1] benefits both providers and customers through its positive effect on various service outcomes. Importantly, communication can address a social dilemma that permeates many congestion-prone social service systems with heterogeneous customer service needs:

[1] In the examples above, it is not clear precisely what "long" and "almost" mean.

the collective overuse of critical resources when utility-maximizing customers fail to internalize the negative externalities they impose on others (Chen and Hasenbein 2020), which creates additional fairness issues when low-need customers overcrowd the system at the expense of high-need customers.

Because price discrimination and centralized admission are difficult to implement in social service settings (e.g., social housing, food banks, emergency departments), in part because they introduce fairness issues of their own, information design is an attractive managerial lever to control demand for scarce resources. Aimed at persuading low-need customers to forego service for the benefit of customers with higher service needs, information design has theoretical appeal as a low-cost non-intrusive mechanism, but it presents service providers with a non-trivial challenge. Because the provision of full information about the operational system represents a challenging task in practice (Anunrojwong et al. 2022) and would create the conditions that result in resource overuse in the first place (Naor 1969), service providers are left with the task of carefully curating information. Complicating things further, for such partial information to have the desired effect on customer behaviour, customers need to find it credible. While standard economic theory suggests the need for a commitment device to render information credible in strategic settings between parties with conflicting incentives, the trouble is that such devices might not be available in practice.

Whether effective (i.e., persuasive) communication can arise in equilibrium between the service provider and their customers, and what the role of commitment is in this interaction, are questions that our study seeks to answer empirically. This is challenging, as solid answers require some control of the environment (e.g., customer needs, ability to commit), the observation of (and variation in) provider-level choices, and the basis to construct a reasonable theoretical benchmark to assess the degree to which communication is effective. Because these conditions are rarely met in field settings, we instead develop a queuing communication game that provides sharp theoretical predictions which we subsequently test under controlled laboratory conditions.

We study the provider and customer decision-making in a setting that shares several of the operational features that characterize service systems in practice, and render effective communication both difficult and important. Arriving in random order to a system, customers decide whether to join a wait list for service, without the ability to actually observe the wait list. Customers are heterogeneous in their need for service, which creates a known dilemma that characterizes many systems with negative externalities (Haviv and Oz 2018): rational and self-interested *low-need* customers join at rates that hurt *high-need* customers. The provider can communicate information to influence customer choices. Specifically, the provider selects a threshold such that arriving customers receive a *short wait* signal if the wait list is shorter than the threshold, and a *long wait* signal otherwise.

Our theoretical analyses show that the effectiveness of such communication depends crucially on the credibility of signals. With commitment, the provider implements the threshold it publicly communicates, which renders signals credible by design. As a result, theory predicts welfare improvements that arise from the provider's ability to persuade low-need customers to balk when it is in their best individual interest to join. In the absence of a commitment device, however, the provider has the incentive to implement a lower threshold than communicated, i.e., to send *inflated long wait* signals to sway away low-need customers from joining the wait list. Realizing this, low-need customers ignore signals to the point where communication effectively breaks down, with no positive welfare effects over the case where the provider does not communicate any information at all.

We test these predictions under controlled laboratory conditions. Our experimental design varies the (human) provider's ability to communicate information about the system states to their (human) customers. We observe that, although a system without any signalling whatsoever fares substantially better than theoretically predicted, communication improves social welfare by mitigating the over-joining behaviour of low-need customers. Our main results concern the role of commitment and its combined effect on the credibility of communication and customer behaviour. We find that the service provider's ability to publicly commit to an information policy influences how both customers and the provider make decisions. However, we observe that the achieved social welfare with and without commitment are statistically indistinguishable. This result stands in sharp contrast to theoretical predictions, but the explanation is subtle: Without commitment, service providers do misreport. But, even though customers can easily detect untruthful communication in our experimental setting, we find that customers still react to signals (although less strongly than in the case with commitment). This allows providers to retain some influence over customers' decisions, and implement policies that send significantly more $l$ong wait signals (in comparison to the case with commitment), which results in the observed welfare equivalence.

The practical implication of our results is that providers can leverage communication as a (potentially) low-cost mechanism to influence customer behaviour towards socially beneficial outcomes in settings with incentive misalignment even when lacking the ability to commit. This is an important insight in light of the fact that credible commitment mechanisms are hard or impossible to design, and it is typically difficult for customers to assess a provider's credibility, in most practical settings.

The remainder of the paper is structured as follows. We review the relevant literature in §2, develop the theoretical model and key predictions in §3, describe our experiments and results in §4, and present in §5 a behavioural model that allows for noisy decision making and accommodates the main behavioural patterns in the data. Finally, in §6, we draw conclusions.

2. Related Literature

Our study relates and contributes to a rich body of theoretical work that studies how information-sharing, via its impact on customer choices, allows service providers to better match demand with scarce supply. One main insight from this theoretical literature is that (rational) customers react to the information environment (Naor 1969, Edelson and Hilderbrand 1975), but that neither fully revealing nor fully concealing information to improve system-level outcomes is uniformly preferable for the entire set of operational and economic parameters (Hassin 1986, Chen and Frank 2004, Shone et al. 2013, Chen and Hasenbein 2020). This has turned information-sharing into a highly studied managerial lever aimed at influencing customer decisions, and the literature has explored various *partial* information structures with different granularities and communication processes; see Ibrahim (2018) and Economou (2021) for comprehensive surveys.

More recent research shows, based on the notion of *Bayesian persuasion* (Kamenica and Gentzkow 2011, Bergemann and Morris 2019), that through intentionally *obfuscating* information, the provider can *persuade* (rational) customers to take desirable actions (Lingenbrink and Iyer 2019, Che and Tercieux 2021). Closest to our setting is the theoretical work of Anunrojwong et al. (2022) which studies information design as a lever to mitigate a dilemma often observed in social service systems: the inefficient and unfair use of critical resources due to customers who fail to internalize the negative externalities that their actions impose on others. In particular, the authors characterize the conditions under which the service provider optimally *persuades* customers with low service needs to forgo service for the benefit of high need customers. We develop a model that shares with Anunrojwong et al. (2022) several features which characterize social service systems, as well as some of its key predictions, but our model is designed specifically for experimental testing.[2]

As theoretical benefits of information-sharing rest on the assumption that customers take information at face value, our study addresses an important overarching issue, namely, whether information is *credible* in the first place. Assuming that credibility holds is problematic in strategic settings where the provider has the incentive to misreport information. Indeed, if incentives between customers and the provider are *sufficiently misaligned*, it is well-known that information-sharing devolves into cheap talk where customer behaviour cannot be influenced (Allon et al. 2011). While many works are silent about the credibility issue, others establish credibility explicitly by assuming that the provider can *commit* to the information-sharing policy; for example, in the persuasion literature (Bergemann and Morris 2019). Whether the theoretical benefits of information-sharing materialize in an equilibrium between providers and customers (who may fall short of the standard rationality assumptions), and whether model assumptions about providers' ability to credibly communicate are justified, are empirical questions that we seek to provide first answers to in this study.

Our focus on the credibility of operational information also relates to a broad and growing empirical literature that documents the (mostly positive) effects of information sharing on customers' perceived value and willingness-to-pay (Buell and Norton 2011, Buell et al. 2017), trust and engagement in

[2] See Kremer and Debo (2016) and Kremer and de Vericourt (2023) for similar examples and the related discussion in Allon and Kremer (2019).

co-production settings (Buell et al. 2021), ex-post perception of wait times and service evaluation (Hui and Tse 1996, Antonides and van Aalst 2002, Munichor and Rafaeli 2007, Ansari et al. 2022), ex-ante beliefs about wait times (Yu et al. 2017), and customer choices such as time spent at service (Webb et al. 2020), service-provider selection (Dong et al. 2019), and queue abandonment (Munichor and Rafaeli 2007, Akşin et al. 2017, Yu et al. 2022).

Our work contributes to this literature for a number of reasons related to the specific characteristics of the communication environment, which we control and manipulate experimentally. First, while this literature indicates that customers may treat information as credible, our work provides a clean test, via the control for strategic (mis)alignment (between the provider and customers) and the manipulation of commitment. Second, to shed light on the mechanisms behind the communication process, our work experimentally studies information design as a game between a service provider and its customers; this requires variation in the provider's selection of information policies, which is difficult to observe in practice. Finally, we provide experimental evidence that quantifies the impact of shared information on individual customers and the system as a whole; this requires a rigorous decision-theoretic benchmark, which is hard to establish in the field.

Finally, while our general focus is on social service systems, we believe that our main insights are relevant to other operational settings that entail strategic communication and where the credibility of information is key, such as supply chain management (Cachon and Lariviere 2001, Anand and Goyal 2009, Özer et al. 2011), retail operations (Allon and Bassamboo 2011, Drakopoulos et al. 2021), or inventory management (Debo and Van Ryzin 2009, Allon et al. 2012, Yu et al. 2015, Schmidt et al. 2015, Aydinliyim et al. 2017, Cui and Shin 2018).

3. Model

We study if and when effective communication can arise between service providers and their customers in social service systems, which are prone to the overuse of resources. To do so, we develop a stylized queuing game that captures the salient features of service systems in practice. First, our model's *system dynamics* are such that customer arrivals and choices dynamically drive the queue length, which captures the state of the system. Second, customers exert *negative externalities* since joining decisions increase wait-related costs for later arrivals. Finally, service providers have an *informational advantage* over customers regarding the system state, which allows them to influence customers through the communication of carefully curated information, e.g., by persuading low-need customers to balk, thereby alleviating congestion for high-need customers.

3.1. Basic Setup: A Wait List Model

There are $\Lambda \in \mathbb{N}$ customers that arrive sequentially, according to a randomly assigned index $k \in \{1, \cdots, \Lambda\}$. Customers do not know their indices in the sequence upon arrival, nor can they observe the indices of other customers (but they know all other system parameters). Customers are of type high-need with probability p_h and of type low-need with probability $1 - p_h$. Unlike low-need customers, high-need customers do not have an outside option, i.e., they always join the system. While the service provider cannot observe the types of individual customers, she knows the probability distribution of customer types. Upon arrival to the system, customers receive some delay-related information, and decide whether to join or balk a wait list to receive service. There is no customer reneging from the wait list. For experimental amenability and without loss of generality, we assume that once all customers have made their join/balk decisions, the (single) service provider commences to serve customers in increasing order of position in the wait list. Service times are deterministic and normalized to 1.

We let U_k denote the utility accrued by customer k, and assume that customers make decisions to maximize their expected utility, $\mathbb{E}[U_k]$. In particular, the (realized) utility that a low-need customer accrues from joining a wait list at the $(q+1)^{st}$ position, i.e., with q customers ahead, is equal to $u_k = r - c(q+1)$. That is, customers earn a reward r upon completion of their service, and incur a delay cost $c(q+1)$, where c represents customer delay sensitivity. The utility that low-need customers experience if they balk is normalized to 0. High-need customers do not have an outside option so that their utility from balking is equal to $-\infty$ (Anunrojwong et al. 2022). We assume that $r - c > 0$ since, otherwise, low-need customers would never join even if the system is empty. We also assume that

$r - c\Lambda < 0$ since, otherwise, low-need customers would never balk even if the system is full. Finally, in our game, the service provider aims at maximizing the *expected social welfare*, $\Omega = \mathbb{E}\left[\sum_{k=1}^{\Lambda} U_k\right]$.

Let $\alpha_w \in [0,1]$ denote a fixed joining probability for low-need customers. We can show that:

PROPOSITION 1. *There is a unique joining probability* $\alpha_w^* \in [0,1)$ *which maximizes the expected social welfare* Ω. *If* $\frac{r}{c} \leq 1 + p_h(\Lambda - 1)$, *then* $\alpha_w^* = 0$. *Otherwise,* $\alpha_w^* = \frac{r - c(1+p_h(\Lambda-1))}{c(\Lambda-1)(1-p_h)} < 1$.

We will see that without delay-information, in equilibrium, low-need customers over-join the system, i.e., they join with a probability that is larger than α_w^*.

3.2. No Communication and Welfare Loss

We start by studying the case in which customers do not receive any information about the wait list. To maximize their expected utility, $\mathbb{E}[U_k]$, low-need customers decide to join whenever $r - c(\mathbb{E}[Q]+1) > 0$, where Q represents the random number of customers in the wait list that a customer encounters upon arrival. To make a joining decision, customers compute $\mathbb{E}[Q]$ based on system parameters and on their belief on how other customers behave. Let $\alpha \in [0,1]$ be the joining probability of low-need customers. Proposition 2 describes the resulting equilibrium in this case.

PROPOSITION 2. *In the case of no communication, a unique equilibrium* $\alpha^* \in [0,1]$ *exists where low-need customers join with probability:*

$$\alpha^* = \begin{cases} 0 & \text{if } \frac{r}{c} \leq \frac{(\Lambda-1)p_h}{2} + 1, \\ \frac{2r - c(2+p_h(\Lambda-1))}{c(\Lambda-1)(1-p_h)} & \text{if } \frac{(\Lambda-1)p_h}{2} + 1 < \frac{r}{c} < \frac{(\Lambda-1)}{2} + 1, \\ 1 & \text{if } \frac{r}{c} \geq \frac{\Lambda-1}{2} + 1. \end{cases}$$

Combining Propositions 1 and 2, we can derive the following result.

COROLLARY 1. *If* $\frac{r}{c} \leq \frac{(\Lambda-1)p_h}{2} + 1$, *then* $\alpha^* = \alpha_w^* = 0$. *Otherwise,* $\alpha^* > \alpha_w^*$.

From Corollary 1, when low-need customers do not join ($\alpha^* = 0$), we know that such behaviour is *socially optimal*, i.e., the service provider does not have an incentive to communicate information to customers to modify their behaviour. However, whenever low-need type customers do join with positive probability ($\alpha^* > 0$), Corollary 1 shows that low-need customers over-join the system. Customers over-join because they disregard the negative externalities that their joining decisions inflict on future arrivals (Haviv and Oz 2018).

In what follows, we restrict attention to the case where $\alpha^* = 1$ and $\alpha_w^* = 0$, i.e., the over-joining of low-need customers is maximal. For this, we assume that $\frac{r}{c} \geq \frac{\Lambda-1}{2} + 1$ and $\frac{r}{c} \leq \Lambda p_h$ (which is satisfied whenever $p_h \geq \frac{\Lambda+1}{2\Lambda}$), where the latter restriction ensures that $\alpha_w^* = 0$, and that the interests of the service provider and low-need customers are sufficiently misaligned.[3] This allows us to derive sharp theoretical predictions to (later) experimentally test the effects of communication, with and without commitment, on customer decision-making and social welfare.

3.3. Communication: Signals, Policy, and Commitment

We next consider the case in which the service provider sends a binary signal $\varsigma = \{s, l\}$, e.g., *s*hort or *l*ong wait, to each arriving customer. In particular, the service provider implements a signalling mechanism with threshold θ such that a customer receives a *s*hort wait signal when arriving to a system with $q < \theta$ customers, and *l*ong wait signal otherwise. We note that optimal signalling mechanisms with a similar threshold structure have been identified in the $M/M/1$ model (Allon et al. 2011, Anunrojwong et al. 2022). The threshold structure is also easy to understand, which is essential for experimental amenability. Although customers do not observe the system state q, they can compute the expected wait list length based on the provider's signals and on their belief about the strategies of others. Without additional information about the service provider's information-sharing policy, i.e, the value of θ, customers must infer the "meaning" of the *l*ong and *s*hort wait messages that they receive. Let $\alpha_s \doteq \mathbb{P}(\text{Join}|\varsigma = s)$ and $\alpha_l \doteq \mathbb{P}(\text{Join}|\varsigma = l)$ be the conditional joining probabilities of low-need customers, given the signal ς. Lemma 1 characterizes customers' best response to a given fixed θ, and the provider's best response to a given fixed customer joining behaviour (α_s, α_l).

[3] In Appendix A.4, we show that for $p_h \geq \frac{r}{c\Lambda}$, communication breaks down in the absence of commitment (cf. Proposition 3). More generally, informative communication can arise without commitment if incentives are sufficiently aligned (Allon et al. 2011).

LEMMA 1. ***(Best Responses)***

(a) Customers: For a given belief about threshold $\theta \in \{0, \cdots, \Lambda\}$*, a unique customer best response* (α_s^*, α_l^*) *exists with* $\alpha_s^* = 1$ *and*

$$\alpha_l^* = \begin{cases} 1 & \text{if } \theta \leq \bar{\bar{\theta}}, \\ \frac{2r - c(2(\theta+1) + p_h(\Lambda - \theta - 1))}{c(1-p_h)(\Lambda - \theta - 1)} & \text{if } \bar{\bar{\theta}} < \theta < \bar{\theta}, \\ 0 & \text{if } \theta \geq \bar{\theta}, \end{cases}$$

where $\bar{\bar{\theta}} = \lfloor \frac{2r - c(\Lambda+1)}{c} \rfloor$, $\bar{\theta} = \lceil \frac{2r - 2c - cp_h(\Lambda-1)}{c(2-p_h)} \rceil$*, and* $0 \leq \bar{\bar{\theta}} < \bar{\theta}$.

(b) Provider: For a given belief about customer joining behaviour $(\alpha_s = 1, \alpha_l \in [0,1])$*, a provider's best response* θ^* *exists with* $\theta^* = 0$ *if* $\alpha_l < 1$*, and* $\theta^* = Any$ *if* $\alpha_l = 1$.

Figure 1 illustrates Lemma 1 and its effect on social welfare, using the parameters from our experiments. For $\theta = 0$, customers would receive only *l*ong signals, rendering such signals uninformative of system states so that low-need customers always join (Lemma 1a, $\alpha_l^* =$ ❶ in the figure). The same argument is true for $\theta = \Lambda$ where only *s*hort signals are sent ($\alpha_s^* =$ ❶ in the figure). As θ increases, customers would receive short signals for an increasing set of system states. As a result, because long signals become strong indicators for long(er) waits, low-need customers are less likely to join when they receive such signals, and never join for large enough $\theta \geq \bar{\theta}$ ($\alpha_l^* =$ ⓿). Notably, while any $\theta > \bar{\theta}$ retains the desired balking behaviour for *l*ong wait messages ($\alpha_l^* =$ ⓿), it would result in more *s*hort wait signals under which low-need customers join ($\alpha_s^* = 1$), thus reducing social welfare.

Indeed, for any given customer joining probability $\alpha_l^* < 1$, Figure 1 illustrates that providers can increase social welfare by setting a threshold $\theta = 0$ (Lemma 1b). This represents a conflict of interest since, as long as low-need customers balk when receiving *l*ong signals, i.e., $\alpha_l < 1$, the provider has the incentive to send only long signals ($\theta = 0$) to dissuade joining. As a result, Proposition 3 characterizes an *uninformative equilibrium* (Allon et al. 2011), where customers entirely disregard signals, and behave as if they were in a setting with no signalling at all, i.e., they always join.

PROPOSITION 3. ***(Ineffective Signals)*** *Equilibria exist such that the provider implements* $\theta^* \in \{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$*, and customers join with probabilities* $\alpha_s^* = 1$ *and* $\alpha_l^* = 1$.

Figure 1 Social Welfare, Communication, and Customer Behaviour ($\Lambda = 8$, $r = 185$, $c = 40$, $p_h = 0.65$)

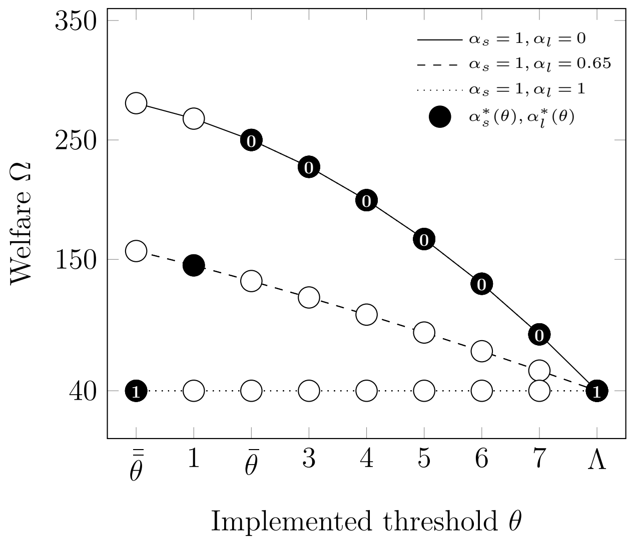

Signals are ineffective to influence customer behaviour when customers have to form beliefs about the information policy θ that generates such signals. But, what if the service provider could communicate the information-sharing policy to customers? We now consider the case where the provider publicly *communicates* a threshold θ' to customers. Whether this communication can effectively mitigate the over-joining behaviour of low-need customers crucially depends on the service provider's credibility. *With commitment*, the service provider makes binding claims θ' about the implemented information-sharing policy θ, i.e., customers are guaranteed that $\theta' = \theta$ so that the provider's communication is credible. In contrast, *without commitment*, the service provider makes non-binding claims θ', i.e., customers know that θ' is not necessarily equal to θ. For communication to be credible in the absence of commitment, it is necessary that the provider does not have the incentive to communicate a threshold θ' that differs from the implemented threshold θ which actually generates the signals. Formally, for a given communicated threshold θ', and customers associated responses $(\alpha_s^*(\theta'), \alpha_l^*(\theta'))$, θ' is credible if, and only if, $\Omega(\theta', \alpha_s^*(\theta'), \alpha_l^*(\theta')) \geq \Omega(\theta, \alpha_s^*(\theta'), \alpha_l^*(\theta'))$ for all possible θ. The following proposition summarizes the relevant equilibria.

Proposition 4. ***(Commitment and No Commitment)*** *(a) With commitment: Let* $g(\bar{\theta}) \doteq 2(r - c\bar{\theta}) + (\Lambda - \bar{\theta})p_h(2r - c(2(1+\bar{\theta}) + p_h(\Lambda - \bar{\theta} - 1)))$. *A unique equilibrium exists between the service provider and customers. If* $g(\bar{\theta}) \geq 0$, *then* $\theta^* = \bar{\theta} < \lfloor \frac{r}{c} \rfloor$, $\alpha_s^* = 1$, *and* $\alpha_l^* = 0$. *Otherwise,* $\theta^* = \bar{\theta} - 1$, $\alpha_s^* = 1$, *and* $\alpha_l^* \in (0,1)$.

(b) Without commitment: Equilibria between the service provider and customers exist such that θ'^* *is any threshold chosen randomly,* θ^* *is a threshold in* $\{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$, $\alpha_s^* = 1$, *and* $\alpha_l^* = 1$.

With commitment, the conflict of interest between the service provider and customers is no longer relevant for customer decision making. This, in turn, allows the provider to select a threshold that anticipates customers' joining best responses as in Lemma 1a. Accordingly, Proposition 4a shows that the provider is able to reduce joining rates with the use of *l*ong wait signals ($\alpha_l^* < 1$). We note that the literature commonly characterizes an *obedient equilibrium* where signals represent action recommendations that are followed with probability 1 (Anunrojwong et al. 2022, Lingenbrink and Iyer 2019). This is without loss of generality if all potential signalling mechanisms are available to the information designer (Bergemann and Morris 2019). In contrast, for the sake of experimental amenability, we consider here a fixed threshold signalling mechanism, under which the obedient equilibrium ($\alpha_s^* = 1$ and $\alpha_l^* = 0$) is not necessarily without loss of generality, as observed in Proposition 4a. In the case of no commitment, Proposition 4b shows that even though the implemented policy is communicated, signals remain ineffective; cf. Proposition 3. This is because the provider has the incentive to implement a lower-than-communicated threshold, i.e., $\theta < \theta'$, rendering communicated thresholds not credible ($\theta' \in \{\bar{\bar{\theta}} + 1, .., \Lambda - 1\}$), or credible but uninformative ($\theta' \in \{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$).

3.4. Informativeness and Persuasion

Since the service provider has an informational advantage over customers, it is natural to ask how much information is transmitted to customers, and how this information shapes the game equilibrium characterized in Proposition 4. Similar to communication games in the economics literature (Fréchette et al. 2022), as a proxy for informativeness, we focus here on how well customers match their actions with true system states. For this, we note that joining customers experience non-negative utility $r - c(q+1) \geq 0$ if, and only if, $q + 1 \leq \frac{r}{c}$, i.e., whenever their position in the wait list does not exceed

$q^* = \lfloor \frac{r}{c} \rfloor$ (Naor 1969). It follows that a low-need customer's decision-making is subsumed to a *binary guess*: to maximize their individual expected utility, customers only need to "guess" correctly if the system state, q, is below q^*. We use the bookmaker informativeness $\mathcal{I}$ metric[4] to quantify the degree of information transmission in our setting (Chicco et al. 2021):

$$\mathcal{I} = \mathbb{P}(\text{Join}|Q < q^*) + \mathbb{P}(\text{Balk}|Q \geq q^*) - 1. \tag{1}$$

From the customer's perspective, $\mathcal{I} = 0$ corresponds to not having any information, and $\mathcal{I} = 1$ corresponds to having all the necessary information to match actions perfectly with the system state. The following Corollary 2 characterizes the informativeness in our game.

Corollary 2. *In equilibrium (Proposition 4):*

(a) With commitment, we have that $\mathbb{P}(\textit{Join}|Q < q^*) \in (0,1)$, $\mathbb{P}(\textit{Balk}|Q \geq q^*) \in (0,1]$, *and* $\mathcal{I} \in (0,1)$.

(b) Without commitment, we have that $\mathbb{P}(\textit{Join}|Q < q^*) = 1$, $\mathbb{P}(\textit{Balk}|Q \geq q^*) = 0$, *and* $\mathcal{I} = 0$.

With commitment, the service provider can achieve $\mathcal{I} = 1$ by implementing a threshold $\theta = q^*$, under which signals perfectly reveal to customers if the wait list is below q^* or not. While the service provider thus has the ability to share all the necessary information for customers to make individually optimal choices, Corollary 2a shows that the service provider finds it optimal to share information only *partially*, i.e., $\mathcal{I} \in (0,1)$. In particular, in line with the notion of *bayesian persuasion* (Kamenica and Gentzkow 2011), a provider trying to encourage balking selects a threshold $\theta^* < q^*$ that *obfuscates* queue states. By pooling realizations of the favorable ($q < q^*$) and unfavorable states ($q \geq q^*$) in the *l*ong wait signal, the provider is able to *persuade* some low-need customers to forego service for the benefit of high-need customers. Formally, we measure the *persuasiveness* of signals as

$$\mathcal{P} = \mathbb{P}(\text{Balk}|\varsigma = l, \theta < q^*). \tag{2}$$

[4] In communication games, the correlation between actions and true states is a standard measure of informativeness (Fréchette et al. 2022). In our case, customers can join with probability 1 in equilibrium (such that the variance is 0), rendering the correlation *undefined*. $\mathcal{I}$ circumvents this problem, and it is proportional to correlation.

It is easy to see from Proposition 4 that $\mathcal{P} \in (0,1]$ in the case of commitment, since customers balk with positive probability when receiving a $l$ong signal under the equilibrium threshold $\theta^* < q^*$. Note that persuasion does not occur for thresholds $\theta \geq q^*$: Although the corresponding long-wait signals would lower joining rates over the case of no signalling, customers in such states need no persuasion as they would balk even if they knew the state q. Finally, for the case without commitment, Corollary 2b shows that $\mathcal{I} = 0$ and $\mathcal{P} = 0$, because customers disregard all information shared by the service provider and join with a fixed probability irrespective of the queue state.

4. Experiment

Our theoretical results show that a service provider can improve social welfare with a properly designed information-sharing policy *if, and only if,* the service provider can credibly commit to it. We designed an experimental study to test these predictions under controlled laboratory conditions that give theory its best shot.

4.1. Task

Participants in our experiment faced the task described in §3.1, in the role of either a service provider or a customer. In each of a total of $T = 40$ experimental rounds, the service provider serves a market of $\Lambda = 8$ potential customers. At the beginning of each round, all customers are randomly and independently assigned a type: high-need with probability $p_h = 0.65$, and low-need with $1 - p_h = 0.35$.

Decisions and system dynamics. Customers arrive to the market sequentially, according to randomly assigned and unknown indices $k \in \{1, \cdots, \Lambda\}$. High-need customers automatically join the wait list, and low-need customers decide whether to join the wait list or not. When a customer joins, the wait list increases by one. After the arrivals and decisions of all Λ customers, the provider delivers service to the customers in the order in which they joined the wait list.

Financials. Customers who join the wait list at position $q + 1$ earn $r - c(q + 1)$, and we set the service value of $r = \$185$ and the delay cost of $c = \$40$. Customers who do not join receive value of \$0 from their outside option. The service provider's objective is to maximize social welfare, which we set as the average utility over all Λ customers.

4.2. Design and Hypotheses

We implement three experimental treatments that vary the service provider's ability to share information to customers and to commit to an information-sharing strategy. In our baseline *NoSignal* treatment, the service provider cannot send to the arriving customers any signals about the length of the wait list. We then implement two treatments in which the service provider can signal information, but that differ in the provider's ability to commit. In the *Commit* treatment, at the beginning of each round, the service provider defines a threshold θ used to generate a binary signal $\varsigma = \{s, l\}$ that customers receive upon their arrival. This threshold is publicly communicated to all customers. Specifically, customers receive the signal s when they arrive to a wait list with less than θ customers, and receive l otherwise. In contrast, service providers in the *NoCommit* treatment first define a threshold θ that is used to generate signals, then publicly communicate a threshold θ'. While customers do not observe θ itself, it is public knowledge that the service provider has the discretion to implement a threshold θ that is different from θ'.

Table 1 summarizes our experimental design and theoretical predictions (for providers' signalling strategies, customers' joining strategies, and welfare) that allow us to test our hypotheses:

HYPOTHESIS 1A. Signals with commitment improve welfare, i.e., $\Omega^{Commit} > \Omega^{NoSignal}$.

HYPOTHESIS 1B. Signals without commitment do not improve welfare, i.e., $\Omega^{NoCommit} = \Omega^{NoSignal}$.

Table 1 Treatments, Sample Sizes, and Theory Predictions.

Treatment	Sessions	N	θ^*	θ'^*	Customer Equil.	Ω	$\mathbb{P}$(Join)	$\mathbb{P}(\text{Join}\mid Q < q^*)$	$\mathbb{P}(\text{Balk}\mid Q \geq q^*)$	$\mathcal{I}$	$\mathcal{P}$
Commit	5	90	2	-	$(\alpha_s^* = 1, \alpha_l^* = 0)$	$250	0.25	0.4	1	0.4	1
NoCommit	7	126	$\{0, 8\}$	Any	$(\alpha_s^* = 1, \alpha_l^* = 1)$	$40	1	1	0	0	0
NoSignal	3	54	-	-	$(\alpha^* = 1)$	$40	1	1	0	0	-

Table 1 also summarizes the mechanisms underlying the hypotheses. Our analyses predict that the ability to commit to a signalling threshold allows the service provider to influence customer behaviour through *informative* signals. Without commitment power, there is no transmission of information, i.e., $\mathcal{I} = 0$, and thus the provider is not able to influence customer behaviour. Indeed, the *NoCommit* and *NoSignal* treatments present the same predictions in terms of $\mathcal{I}$ and customer behaviour. In contrast, with commitment power, signals are *partially informative*, i.e., $\mathcal{I} \in (0,1)$, such that the service provider persuades low-need customers to balk: $\mathcal{P} = 1$.

4.3. Discussion of Design Choices

We now briefly discuss our main experimental design choices.

Human service providers. Because we want to study if effective communication can arise in equilibrium, our experiments feature both human customers and human service providers. Because the presence of a human service provider might affect customer decisions for behavioural reasons not considered in our theoretical analyses, and to allow for a clean comparison with the communication treatments (*Commit, NoCommit*), we include a human service provider even in the *NoSignal* treatment where service providers do not make decisions that affect the game outcomes.

Strategy method. We elicit from participants the choices, join or balk, that they would make as low-need customers before they learn about their actual type for the round (Figure 2b). That is, given communicated thresholds θ', in treatments *Commit* and *NoCommit*, we elicit from participants their choices in response to signals that they could receive. The strategy method, extensively used in the experimental economics literature (see e.g. Brandts 2011, Beer et al. 2022), addresses two challenges related to data availability and consistency between model setting and experimental implementation. On the issue of data, the strategy method allows us to observe participants' choices as low-need customers even in rounds that assigns them a high-need type, and to fully understand customer behaviour even in scenarios that are unlikely to happen in our data. Additionally, the strategy method is crucial to establish consistency between our experimental implementation and the unobservable rank assumption that is at the core of our theoretical developments. We need to ensure that participants cannot, even imperfectly, infer their rank k, e.g., from how long they have waited after all participants moved to the next round. The strategy method rules out such inferences.

Feedback and learning. At the end of each round, using the strategies elicited from participants (service provider communication decisions and customer joining decisions) and the realizations of

Figure 2 Screenshots (treatment *NoCommit*)

(a) Service provider: Elicit θ and θ' (b) Customer: Elicit $a(\varsigma = s|\theta')$ and $a(\varsigma = l|\theta')$

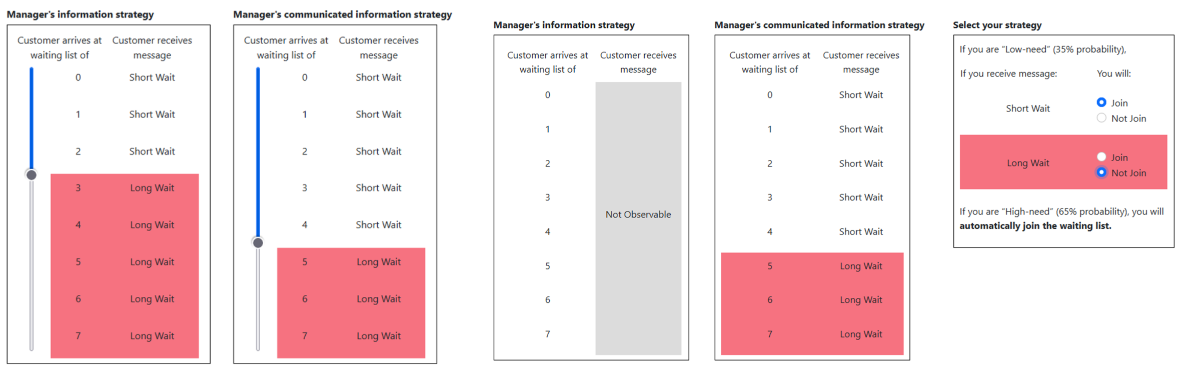

Notes: This figure presents screenshots once service providers and customers have made their selections. To avoid anchoring, players do not see any default selection.

random events (customer type and index), the computer simulates the system; see Figure 3. All participants, regardless of their roles, receive full feedback which includes the communicated threshold, customer order of arrival, customer types, signals received, customer decisions, customer positions in the wait list, and the utilities generated. This ensures that the joining decisions of customers are based solely on their payoff function, and not on other factors such as their desire to gain information for future rounds, e.g., customers could join to know how others behave. Although service systems in practice rarely provide such extensive information, it is desirable to do so in our experimental implementation because it gives theory its best shot, by reinforcing participant understanding about the dynamics in the queue, the signalling threshold, and the computation of their payoffs.

Finally, we note that the full access to results also means that customers in the *NoCommit* treatment would observe whether the provider's implemented threshold matched the communicated threshold. Since there is no guarantee that all participants would arrive at the right conclusion (or that they engage at all in such computation), to control for the influence of the availability of this information, we display the service provider's communicated and implemented policy side-by-side. In §4.7, we discuss the *NoCommit* treatment findings in light of this latter experimental design choice.

Parameters. For our chosen parameter values (r, c, p_h, Λ), all low-need customers should theoretically join the system when no information is communicated (see Proposition 2), thus creating a loss

Figure 3 End-of-round Feedback (treatment *Commit*)

Customer	Order of Arrival	Type	Arriving at waiting list of	Received Message	Decision	Position in Waiting List (P)	Service Value (S)	Waiting Cost (W) = \$40*(P)	Customer Payoff = (S)-(W)
CC	1	High-need	0	Short Wait	Join	1	\$185	\$40	\$145
CF	2	High-need	1	Short Wait	Join	2	\$185	\$80	\$105
CA	**3**	**Low-need**	**2**	**Short Wait**	**Join**	**3**	**\$185**	**\$120**	**\$65**
CH	4	High-need	3	Long Wait	Join	4	\$185	\$160	\$25
CB	5	Low-need	4	Long Wait	Join	5	\$185	\$200	-\$15
CE	6	Low-need	5	Long Wait	Not Join	None	\$0	\$0	\$0
CG	7	Low-need	5	Long Wait	Not Join	None	\$0	\$0	\$0
CD	8	High-need	5	Long Wait	Join	6	\$185	\$240	-\$55
								Manager's Payoff = Average Customer Payoff	\$33.75

Notes: This figure illustrates the results table that customer CA observes.

in social welfare that can be mitigated by a properly designed communication strategy (see §3.3). Our choice of $\Lambda = 8$ deserves further comment. Although service environments typically involve a large number of customers that interact with the system, the practical realities of our laboratory environment required a careful selection of Λ. We selected a Λ that is small enough to be able to increase the number of sessions (i.e., independent observations) as much as possible given a fixed pool of subjects, while large enough to retain the complexity of the decision-making task that arises from the interaction of multiple players.

4.4. Software, Recruitment, and Payment

We implemented the experiment in oTree (Chen et al. 2016). In total, 270 participants were included in our study and each subject participated in one treatment only. We recruited participants from a subject pool associated with the experimental laboratory at a large public university in Europe. In each session, after arriving at the laboratory, participants were randomly assigned to isolated cubicles and read instructions presented on-screen in an easy-to-navigate format. Participants in all treatments then played five practice rounds of the *NoSignal* treatment to familiarize themselves with the task and computer interface with the option to ask clarification questions. The computer then randomly assigned to each participant a role (provider or customer) that they would keep for the duration of the experiment. Exactly 18 participants were in a session, for 2 cohorts of 9 (1 provider and 8 customers) and participants did not know the session size. At the beginning of each round, each service provider was randomly matched with 8 customers to eliminate reputational concerns.

Each session lasted around 90 minutes, and subjects were paid €7 for their participation, in cash at the end of the session, plus a bonus based on their average payoff from all 40 rounds, at a conversion rate of €0.5 for each $1 experimental dollar. Total earnings ranged from €7.44 to €33.38 with an average of €19. Prior to collecting data for our studies, we pre-registered the main research hypotheses, key dependent and independent variables, analysis plans, target sample sizes, and exclusion criteria. The full pre-registration documents are available at `https://aspredicted.org/H92_D94` and `https://aspredicted.org/MXN_MJZ`. [5]

4.5. Data and Analysis

At the most granular level, in each round $t \in \{1, \ldots, T\}$ and for each cohort $c \in \{1, \ldots, C\}$, we observe the provider j's implemented threshold decision θ_{jtc} in *Commit* and *NoCommit*, as well as the communicated threshold decision θ'_{jtc} in *NoCommit*. For each customer $i \in \{1, \cdots, \Lambda\}$, we observe their strategy $A_{itc} \equiv \{a_{itc}\}$ in *NoSignal*, $A_{itc} \equiv \{a_{itc}(\varsigma = s|\theta_{jtc}), a_{itc}(\varsigma = l|\theta_{jtc})\}$ in *Commit*, and $A_{itc} \equiv \{a_{itc}(\varsigma = s|\theta'_{jtc}), a_{itc}(\varsigma = l|\theta'_{jtc})\}$ in *NoCommit*, where $a_{itc} \in \{0, 1\}$ is such that 1 represents the *Join* action and 0 the *Balk* action. We also observe the realized welfare w_{ct}, but we use for our analyses the *expected* welfare Ω_{ct} conditional on the observed joining strategies A_{itc} and implemented thresholds θ_{jtc}. This eliminates from our data the impact of variability from random realizations of customer arrival indices and types, which provides a clean comparison with theoretical predictions. In Appendix B, we describe how we compute Ω_{ct} and, similarly, other key metrics that we analyze.

To compare metrics across treatments, or against theoretical predictions, we use session-level averages as the unit of analysis for our statistical tests (i.e., t-tests). In addition, we also report several regression-based analyses with standard errors clustered at the appropriate level to accommodate the dependency of observations in our data.

[5] While we aimed for 8 sessions in the *NoCommit* treatment, the effective size of the subject pool did not allow for the last session despite extensive recruitment efforts. However, based on our current results, we suspect that such additional session would have most likely only strengthen our claims.

4.6. Results

Table 2 displays the main results for total welfare, service provider behaviour, and customer behaviour. We next discuss these in more detail.

Table 2 Experimental Results (standard deviations in parentheses)

	Welfare	$\mathbb{P}(\text{Join})$	$\mathbb{P}(\text{Join}\|Q < q^*)$	$\mathbb{P}(\text{Balk}\|Q \geq q^*)$	$\mathcal{I}$	$\mathcal{P}$
NoSignals	172.84	0.57	0.57	0.43	0	-
	(22.59)	(0.09)	(0.09)	(0.09)	(0.00)	-
Commit	200.70	0.48	0.66	0.74	0.41	0.74
	(10.86)	(0.05)	(0.06)	(0.09)	(0.12)	(0.09)
NoCommit	198.70	0.48	0.61	0.68	0.29	0.66
	(10.20)	(0.05)	(0.11)	(0.09)	(0.18)	(0.09)

4.6.1. Social Welfare. Figure 4 presents, for each treatment, the average implied social welfare over the course of the experiment.

Figure 4 Social Welfare

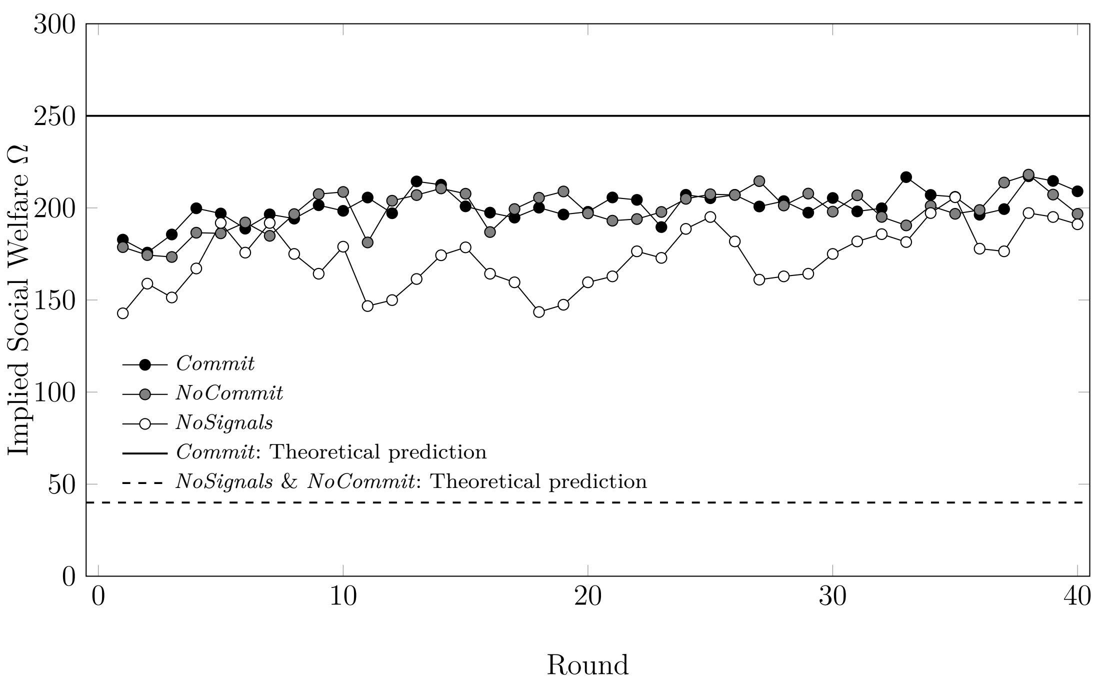

We make several observations. Notably, the welfare in the *NoSignal* treatment is substantially better than theoretically predicted (172.84 vs. 40, $p = 0.005$). Although our data thus leaves less room (than theoretically predicted) for effective communication to increase social welfare, we observe that *Commit* improves welfare over *NoSignal* (200.7 vs. 172.84, $p = 0.026$), despite falling short of theory predictions (200.7 vs. 250, $p < 0.001$). Overall, the data supports Hypothesis 1A ($\Omega^{Commit} > \Omega^{NoSignal}$). We observe that welfare also improves under *NoCommit* (198.7 vs. 172.84, $p = 0.016$), thus rejecting Hypothesis 1B ($\Omega^{NoCommit} = \Omega^{NoSignal}$) that signals without commitment do not improve social welfare. In fact, surprisingly, social welfare does not differ between the two signal treatments (200.7 vs. 198.7, $p = 0.376$), leading to our first result.

RESULT 1. *Signals with and without commitment improve social welfare to the same extent.*

Table 5 in Appendix C.1 presents OLS regressions that support the aforementioned findings. The aggregate nature of our analysis thus far leaves open the question of why communication is effective even without commitment. We next study the reasons for this result in more detail.

4.6.2. Informativeness and Persuasiveness. A possible reason for the higher than predicted effectiveness of the *NoCommit* treatment is that signals are informative and persuasive even though they theoretically should not be. Recall that, as a proxy for informativeness, the metric $\mathcal{I}$ captures how well customers match their choices with the actual system state. Not surprisingly, we observe that the average $\mathcal{I}$ is higher in *Commit* than in *NoSignal* (0.41 vs. 0, $p = 0.001$). In fact, signal informativeness in *Commit* is about as high as theoretically predicted (0.41 vs. 0.4, $p = 0.908$). Contrary to theoretical predictions, our data further shows that signals in *NoCommit* are also informative even though theory predicts they are not (0.29 vs. 0, $p = 0.003$), and no less so than in *Commit* (0.29 vs. 0.41, $p = 0.127$).

RESULT 2. *Signals with and without commitment transmit the same amount of information.*

To further analyze the welfare implications of this result, recall that service providers have the incentive to convince as many low-need customers to balk as possible, and that they can do so via two related mechanisms, formally encapsulated in our informativeness metric $\mathcal{I} = \mathbb{P}(\text{Join}|Q < q^*) + \mathbb{P}(\text{Balk}|Q \geq q^*) - 1$. First, service providers aim to reduce customers' propensity to join when it is in the customers' best interest to join, $\mathbb{P}(\text{Join}|Q < q^*)$. Table 2 shows that, relative to *NoSignal* (where $\mathbb{P}(\text{Join}|Q < q^*) = 0.57$), service providers are not significantly able to send signals that reduce customers' ability to join, in *Commit* (0.66 vs. 0.57, $p = 0.063$) and in *NoCommit* (0.61 vs. 0.57, $p = 0.306$). Second, service providers have the incentive to increase customers' ability to balk when it is customers' best interest to balk, $\mathbb{P}(\text{Balk}|Q \geq q^*)$. Indeed, relative to *NoSignal* (where $\mathbb{P}(\text{Balk}|Q \geq q^*) = 0.43$), service providers are able to improve customers' ability to balk, in *Commit* (0.74 vs. 0.43, $p = 0.001$) and in *NoCommit* (0.68 vs. 0.43, $p = 0.002$).

Signals are influential, but are they persuasive? Recall that we measure persuasiveness as $\mathcal{P} = \mathbb{P}(\text{Balk}|\varsigma = l)$ for $\theta < q^*$, which captures the propensity of customers to balk when receiving an l signal when the service provider uses a persuasive threshold, i.e., $\theta < q^*$. Table 2 shows that, relative to the overall baseline balking probability in the *NoSignal* treatment (which is on average 0.43), $\mathcal{P}$ is higher in *Commit* (0.74 vs. 0.43, $p = 0.001$) and in *NoCommit* (0.66 vs. 0.43, $p = 0.002$). Indeed, the level of persuasion is the same in both treatments (0.74 vs. 0.66, $p = 0.071$).

RESULT 3. *Signals with and without commitment are equally persuasive.*

Although there is indeed persuasion in the data, the fact that service providers fail to reduce customers' ability to join (i.e., $\mathbb{P}(\text{Join}|Q < q^*)$) suggests that the frequency in which persuasive thresholds are used and the extent to which customers are persuaded is not sufficient, leaving potential social welfare improvements on the table.

4.6.3. Choices: Customers. A possible reason for the observed aggregate-level equivalence in social welfare between *Commit* and *NoCommit* is that customers react in the same way to signals regardless of whether or not the signals are credible (via commitment). We now study if commitment, or lack thereof, has an effect on customers' joining decisions. Figures 5a and 5b present the average customer joining probabilities, for a given communicated threshold and signal, calculated based on customers' strategies A_{itc}. We observe that the joining probability in *Commit* decreases in the communicated threshold θ', more than predicted for *s*hort signals (where theory predicts that choices are not sensitive to θ' at all; Figure 5a), but less than predicted for *l*ong signals (Figure 5b). In contrast, the joining probability in *NoCommit* is relatively more constant across communicated thresholds, and it is generally lower than theory predictions (Figures 5a and 5b). To test these observations formally, we run several logistic regressions with customer joining decisions as dependent variable (Appendix C.2, Table 6). The results show that customers in the *Commit* and *NoCommit* treatments react differently to the communicated thresholds. In particular, for a *s*hort wait signal, the probability to join in the *Commit* treatment decreases significantly in the communicated threshold, but does not vary significantly in the *NoCommit* treatment. For the *l*ong wait signal, the regressions show that the probability to join decreases significantly in the communicated threshold in both *Commit* and *NoCommit* treatments, but decreases more markedly in the *Commit* treatment.

Result 4. *Customers are more sensitive to thresholds with commitment.*

The post-experimental self-reports (Appendix C.4) provide further evidence for this result, showing that customers in the *Commit* treatment felt that the service provider's information strategy had more impact on their decisions in comparison to those in *NoCommit* (5.65 vs. 4.41, $p = 0.005$).

Figure 5 Customer Joining Decisions and Providers' Communicated (θ') and Implemented (θ) Thresholds

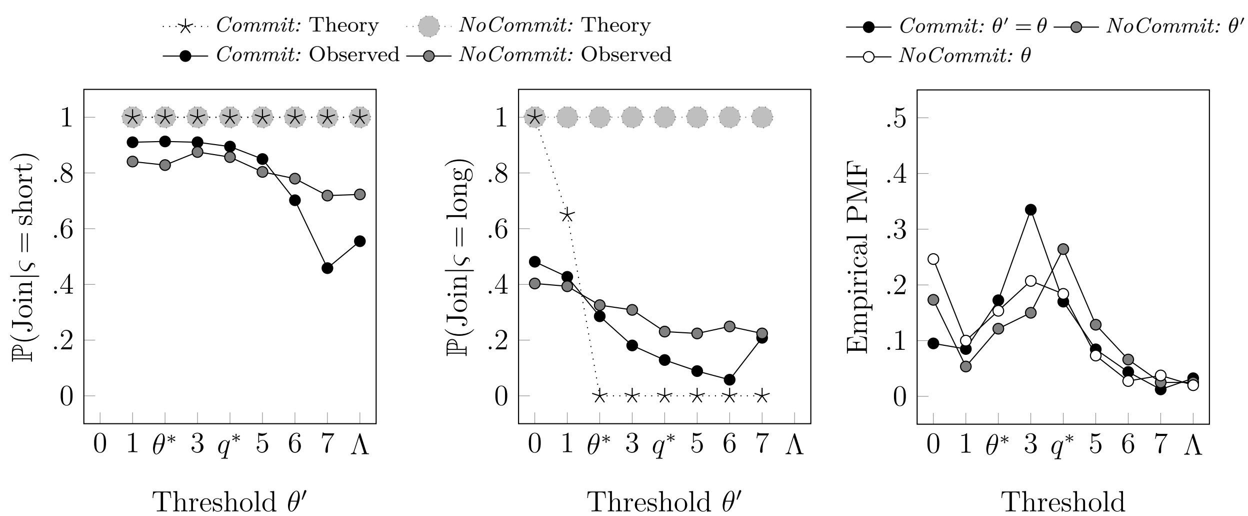

(a) Customers ($\varsigma = \text{short}$) (b) Customers ($\varsigma = \text{long}$) (c) Provider: threshold choices

4.6.4. Choices: Service Providers. Result 4 has profound implications on the decisions made by service providers, and hence for the game equilibrium and resulting system performance. Figure 5c shows service providers' threshold selections. Because customers in *NoCommit* respond to thresholds even when theory predicts they should not, we see that service providers strategically implement and communicate thresholds to influence customer behaviour even in the absence of commitment. However, because customers are less sensitive to thresholds in *NoCommit*, we see that efficient (towards the goal of increasing social welfare) communication requires service providers to implement and communicate thresholds differently when they cannot credibly commit. Indeed, we conduct two-sample Kolmogorov-Smirnov tests and find that the distribution of thresholds in *Commit* is significantly different from both the distribution of implemented thresholds ($D = 0.19$, p-value < 0.001), and communicated thresholds ($D = 0.19$, p-value < 0.001) in *NoCommit*.

Result 5. *Commitment influences how service providers implement and communicate thresholds.*

Importantly, we observe that service providers in *NoCommit* do not communicate truthfully. Specifically, on average, service providers implement a threshold that is different from the one communicated 52% of the time. To quantify the direction and magnitude of this miscommunication, we define the metric $Lie = \theta' - \theta$. Conditional on service providers' decision to lie (i.e., $Lie \neq 0$), implemented thresholds are, on average, about three thresholds apart from what service providers communicate ($|Lie| = 2.97 > 0, p < 0.001$), and the deviations appear to be strategic: the average Lie of 1.42 ($> 0, p = 0.035$) shows that service providers implement thresholds that are systematically lower than the communicated ones. Given that customers do not entirely discard the information they receive (Result 4), service providers indeed have the incentive to implement $\theta < \theta'$ given that their payoff (i.e., social welfare) increases in low-need customers' balking probability (see Figure 1).

Result 6. *Without commitment, service providers implement lower thresholds than communicated.*

Does untruthful communication pay off for service providers, i.e., does it increase social welfare in the *NoCommit* treatment? To formally answer the question, we estimate simple OLS regressions on the data from the *NoCommit* treatment (Table 7, Appendix C.3). For example, we have $\Omega_{ct} = \beta_0 + \beta_1 Lie.Type(\theta' > \theta) + \beta_2 Lie.Type(\theta' > \theta) + \beta_3 Round + \beta_4 Gender.M + \epsilon_c$, where $Lie.Type$ represent indicator variables. The results show that service providers who inflated communicated thresholds such that $\theta' > \theta$ achieved a higher social welfare in comparison to those who were honest, i.e., $\theta' = \theta$.

Moreover, we see that service providers who selected thresholds such that $\theta' < \theta$ achieved a lower social welfare in comparison to those who were honest.

RESULT 7. *Service providers that communicate inflated thresholds achieve higher social welfare.*

Our observation that service providers lie (Result 6), yet successfully (Result 7) manage to influence customers with non-credible signals (Results 2 and 3), deserves further discussion in light of our design that allows customers to observe both implemented threshold θ and communicated threshold θ' at the end of each round (Section 4.3, Figure 2). At the surface, this simple lie detection device effectively increases the lying cost and hence may act like an informal behavioural commitment device - fearing detection and the resulting game collapse to the babbling equilibrium and the low welfare it implies, service providers may refrain from implementing lower thresholds than communicated. While consistent with the observed welfare equivalence (Result 1), such a mechanism is not consistent with the extent of lying we observe. Communication without formal commitment works, but it is not because the high lie detection probability turns service providers into trustworthy communicators.[6]

4.7. Discussion

We observe that both *Commit* and *NoCommit* increase welfare over the *NoSignal* treatment (Result 1). In other words, communication can effectively improve social welfare even when the service provider lacks the ability to commit, in contrast to the prediction of standard theory (Section 3.3). Notably, we observe the welfare equivalence between *Commit* and *NoCommit* despite the fact that commitment *does* change both customers' and service providers' behaviours (Results 4 and 6), which points to a subtle interplay not predicted by theory. Because customers observe or expect considerable lying, service providers in *NoCommit* have less influence over customers' balking decisions than in *Commit* (Result 4). Nonetheless, customers are not entirely insensitive to information and service

[6] While speculative, it is not obvious that a design that prevents customers from directly observing each lie would result in the babbling equilibrium predicted by theory. In fact, it might further benefit the *NoCommit* if an increased tendency to implement lower thresholds than communicated coincides with sustained customer trust in the informativeness of signals (not an unlikely scenario given that customers in our data exhibit significant trust in the informativeness of signals despite the fact that they can easily detect the extant lying on the part of the providers).

providers in *NoCommit* implement lower thresholds than communicated (Result 6), slightly increasing the balking probability for customers receiving the *l*ong wait signal (by increasing the communicated threshold) and also increasing the occurrence of *l*ong wait signals in more queue states (by decreasing the implemented thresholds). By the same token, since service providers set lower thresholds in *NoCommit* than in *Commit*, customers receive *l*ong wait signals more frequently in *NoCommit*. Overall, while the lack of commitment limits service providers' influence over customers' decisions, it allows them to lie and implement lower thresholds. This behaviour is intriguing, as it reveals a delicate balance between service providers and customers, resulting in information transmission and improved social welfare even in the absence of a commitment device.

5. Explaining and Modelling Behaviour

Our experimental results depart from theoretical predictions in a number ways at the welfare level. First, we find that a system without communication performs significantly better than predicted, as low-need customers do not overjoin as much as anticipated. Second, although communication with commitment reduces customer overjoining, it falls short from the predicted improvement. Lastly, we find that communication without commitment improves social welfare, and importantly, to the same extent as with commitment. Softening the apparent challenge that these results pose to standard theory, we observe that individual-level decision-making, while noisy and biased, qualitatively aligns with theoretical predictions. For example, in a system without communication, customers join more than they balk. In a system with communication and commitment, customers join more when receiving short signals than when receiving long signals, and their joining when receiving long signals decreases in the communicated threshold. And, while there is variability in providers' choices, they usually pick thresholds that are close to the optimal one. Finally, in a system without commitment, providers implement lower than communicated thresholds, i.e., they lie, and customers are less sensitive to changes in the communicated threshold in comparison to the case with commitment.

The above observations qualitatively align with the notion of a *quantal response*, where all possible choices are candidates for selection, but more attractive alternatives, yielding higher utility, are chosen with larger probability (Chen et al. 2012). Accordingly, to reconcile our observations, we present a structural behavioural model that uses the Quantal Response Equilibrium (QRE) framework (Goeree et al. 2016), which has recently been popularized in the behavioural operations literature (e.g., Su 2008, Chen et al. 2012, Huang et al. 2013, Kremer and Debo 2016, Goldschmidt et al. 2021).The QRE is appealing as it accounts for noise in decision making, allows to incorporate relevant behavioural factors, e.g., lying aversion, and can accommodate the main departures from standard decision theory without leaving the grounds of this theory entirely.

5.1. Customer QRE (*NoSignals*)

The QRE studies decision-making under a notion of *noisy optimization*, where players are not able to *always* identify their best responses against the actions of others. Formally, player i chooses a *noisy* best response, i.e., a distribution over actions, by maximizing $\mathbb{E}[\text{Utility}] + \epsilon_i$ instead of $\mathbb{E}[\text{Utility}]$, where ϵ_i is a noise term. By assuming that the ϵ_i terms are independently and identically distributed according to a mean-zero Gumbel distribution with scale parameter $\beta_i > 0$, we obtain a logit specification for the choice probabilities (McFadden 1981). The parameter β_i captures player i's level of *bounded rationality* (Chen et al. 2012, Huang et al. 2013), as such a parameter is proportional to the standard deviation of the noise term ϵ_i ($\approx 0.78\beta_i$). When $\beta_i \to 0$, player i chooses the utility-maximizing alternative with certainty, i.e., the theoretical prediction under full rationality. At the other extreme, when $\beta_i \to \infty$, player i lacks the ability to make any rational judgement and, thus, randomizes over all alternatives with equal probabilities. For the sake of parsimony, we use a common parameter β for every customer. In (3), we present the logit specification that characterizes the low-need customers' joining probability under QRE:

$$\varphi(\beta) = \frac{e^{(r-c(\mathbb{E}[Q]+1))/\beta}}{1 + e^{(r-c(\mathbb{E}[Q]+1))/\beta}}. \tag{3}$$

We recall that high-need customers do not have an outside option, so they always join.

The QRE framework preserves the notion of *equilibrium*, in the sense that players choose distributions over actions with correct expectations and beliefs about the choices of others. In line with this, in the case of *NoSignals*, the QRE is the solution of the fixed-point problem given by equation (3). We defer the reader to Appendix D for further details on the QRE, and the computation of $\mathbb{E}[Q]$.

5.2. Provider-Customer QRE (*Commit, NoCommit*)

The fundamental logic of the QRE framework is applicable to settings with communication, but requires some modeling adjustments to reflect the additional complexity. Equations (4) and (5), respectively, present the logit specifications for low-need customers' joining probabilities $\varphi_\varsigma(\theta')$ and the provider's threshold joint probability mass function $\vartheta(\theta,\theta')$:

$$\varphi_\varsigma(\theta';\beta,\beta_m,\kappa) = \frac{e^{(r-c(\mathbb{E}[Q|\varsigma,\theta']+1))/\beta}}{1+e^{(r-c(\mathbb{E}[Q|\varsigma=l,\theta']+1))/\beta}} \quad \text{for} \quad \theta'=0,\cdots,\Lambda, \tag{4}$$

$$\vartheta(\theta,\theta';\beta,\beta_m,\kappa) = \frac{e^{(\Omega(\varphi_s(\theta'),\varphi_l(\theta'),\theta)-\kappa|\theta'-\theta|)/\beta_m}}{\sum_\theta \sum_{\theta'} e^{(\Omega(\varphi_s(\theta'),\varphi_l(\theta'),\theta)-\kappa|\theta'-\theta|)\beta_m}} \quad \text{for} \quad \theta,\theta'=0,\cdots,\Lambda. \tag{5}$$

Customer equations (4) expand equation (3) by conditioning on a signal ς and a communicated threshold θ'. In *NoCommit*, since we have a simultaneous game where θ' is not necessarily equal to the actual implemented threshold θ, customers form expectations $\mathbb{E}[Q|\varsigma,\theta']$ based on their beliefs about the joining of others, and their beliefs about the provider's joint probability distribution over the pair of thresholds θ and θ', which is characterized by equation (5). Consistent with the notion of preferences for truth-telling in the behavioural economics literature (Rosenbaum et al. 2014, Abeler et al. 2019), in (5) we capture a psychological *lying cost* with lying aversion parameter κ, which is proportional to the extent to which the provider misreports thresholds, $\kappa|\theta'-\theta|$ (Özer et al. 2011).

In *Commit*, since we have a sequential game where $\theta'=\theta$, customers form expectations $\mathbb{E}[Q|\varsigma,\theta']$ based on a *realized* implemented threshold $\theta=\theta'$, and their belief about the joining of others. In this case, (5) reduces to a univariate probability distribution, where the provider selects implemented thresholds in a noisy fashion. By the same token, because $\theta=\theta'$ (providers are not able to lie), the lying cost term in (5) only applies to *NoCommit*. The QRE represents the solution of the system of equations (4)-(5). We defer the reader to Appendix D for further details on the QRE, and the computation of $\mathbb{E}[Q|\varsigma,\theta']$ and social welfare $\Omega(\varphi_s(\theta'),\varphi_l(\theta'),\theta)$.

5.3. Estimation

For each treatment, we estimate the set of relevant parameters, β,β_m,κ, using standard maximum likelihood techniques (see Appendix D for details). Table 3 summarizes the estimations, and presents implied values of aggregated metrics under such estimations (see Appendix D for details on their computation). Table 3 includes nested versions that impose certain constraints on the set of relevant parameters. Specifically, we estimate models with $\beta = \beta_m$ to assess whether role-specific "decision noise" adds value for *Commit* and *NoCommit*. In the same spirit, we estimate $\kappa = 0$ to assess the added explanatory power of the lying aversion term in *NoCommit*. Based on a series of log-likelihood ratio tests, for both communication treatments, the unconstrained models (2b and 3d) provide the best fit for our data. Figure 6 presents customer and provider's choice predictions under these full models. We discuss the results from these models next.

5.4. Results

At a high level, we see from Table 3 that $\hat{\beta}$ and $\hat{\beta}_m$ are greater in *NoCommit* than in *Commit*, which is consistent with the intuitive notion that *NoCommit* represents a more complex decision-making environment. Interestingly, $\hat{\beta}$ is highest in *NoSignals*, which indicates that providers' signalling (even without commitment) seems to reduce customers' task complexity. We also see in Table 3 that the orders between the implied aggregate metrics across treatments under QRE are consistent with our experimental data (see Table 2). Indeed, the point estimates for social welfare, $\mathbb{P}(\text{Join}|Q < q^*)$, $\mathbb{P}(\text{Balk}|Q \geq q^*)$, $\mathcal{I}$, and $\mathcal{P}$ are all highest in *Commit*, then in *NoCommit*, and the lowest in *NoSignals*. Nevertheless, the actual values are slightly *biased* since our QRE specification *over-estimates* the joining rates of low-need customers observed in the data (see Figures 6a and 6c). Consequently, the implied social welfare under the estimated QRE, while higher in the communication treatments, is not as high as in the data.

In terms of customer and provider choices, we observe that the notion of *quantal response* in the QRE is able to accommodate the main trends in the data. For example, in *Commit*, the observed threshold frequencies in the data are well captured by the QRE predictions since welfare under QRE follows a similar unimodal trend conditional on customers' quantal responses (see Figure 6b). Also we observe in *Commit* an overall decreasing joining probability, for both short and long signals, as communicated thresholds θ' increase (see Figure 6a). The QRE predictions capture this as the expected

Table 3 QRE Estimation Results (standard errors in parentheses) and Implied Aggregate Metrics

	NoSignals	*Commit*		*NoCommit*			
	(1) Full	(2a) $\beta=\beta_m$	(2b) Full	(3a) $\beta=\beta_m,\kappa=0$	(3b) $\beta=\beta_m$	(3c) $\kappa=0$	(3d) Full
$\hat{\beta}$	91.12* (16.63)	38.10* (0.05)	39.38* (0.10)	66.64* (1.07)	62.63* (1.17)	52.63* (1.42)	57.59* (1.44)
$\hat{\beta}_m$	-	-	13.67* (1.05)	-	-	197.12* (32.25)	113.18* (16.99)
$\hat{\kappa}$	-	-	-	-	28.28* (1.44)	-	44.21* (6.03)
Obs. [Customers/Providers]	1,920/-	6,400/400		8,960/560			
Log-Likelihood	-1,311.51	-3,431.52	-3,380.37	-7,294.67	-7,064.32	-7,246.51	-7,052.61
LL-ratio test (vs. Full)	-	$\chi^2(1)=102.29^*$	-	$\chi^2(2)=484.11^*$	$\chi^2(1)=23.41^*$	$\chi^2(1)=387.78^*$	-
Social Welfare	172.84	174.87	179.44	175.04	176.30	168.08	174.19
$\mathbb{P}$(Join)	0.57	0.58	0.57	0.57	0.57	0.60	0.58
$\mathbb{P}(\text{Join}\mid Q<q^*)$	0.57	0.74	0.75	0.64	0.66	0.69	0.67
$\mathbb{P}(\text{Balk}\mid Q\geq q^*)$	0.43	0.59	0.64	0.51	0.53	0.51	0.53
$\mathcal{I}$	0	0.32	0.39	0.16	0.19	0.19	0.21
$\mathcal{P}$	-	0.57	0.61	0.57	0.58	0.61	0.60

$^*p<0.001$.

payoff from joining given a short (long) signal under QRE decreases (increases) as θ' increases. Similarly, in *NoCommit*, we observe that the overall joining behaviour is less responsive towards changes in θ', in comparison to *Commit* (see Figures 6a and 6c). The QRE predictions capture this since, due to providers' lying behaviour, customers are not able to infer θ perfectly under QRE, such that the expected payoff does not change much in θ' (as in *Commit*). At the same time, for customers to infer θ via the communicated θ', and thus for their joining behaviour to not be the same for all θ' under QRE (see Figures 6a and 6c), it is required that providers incur lying costs, i.e., we have $\kappa>0$. Lying costs are also required for providers to select communicated thresholds in a non-uniform fashion under QRE (see Figure 6d). Finally, we note that while models with lying costs fit the data better and explain the aforementioned patterns, bounded rationality is sufficient to account for social welfare improvements over both the theoretical prediction and the no communication baseline (see model 3a in Table 3).

6. Conclusions

Many service systems suffer from a social dilemma that arises when customers with relatively low service needs overuse scarce resources at the expense of customers with higher service needs. Because

Figure 6 Observation vs. QRE Predictions (models 2b and 3d from Table 3)

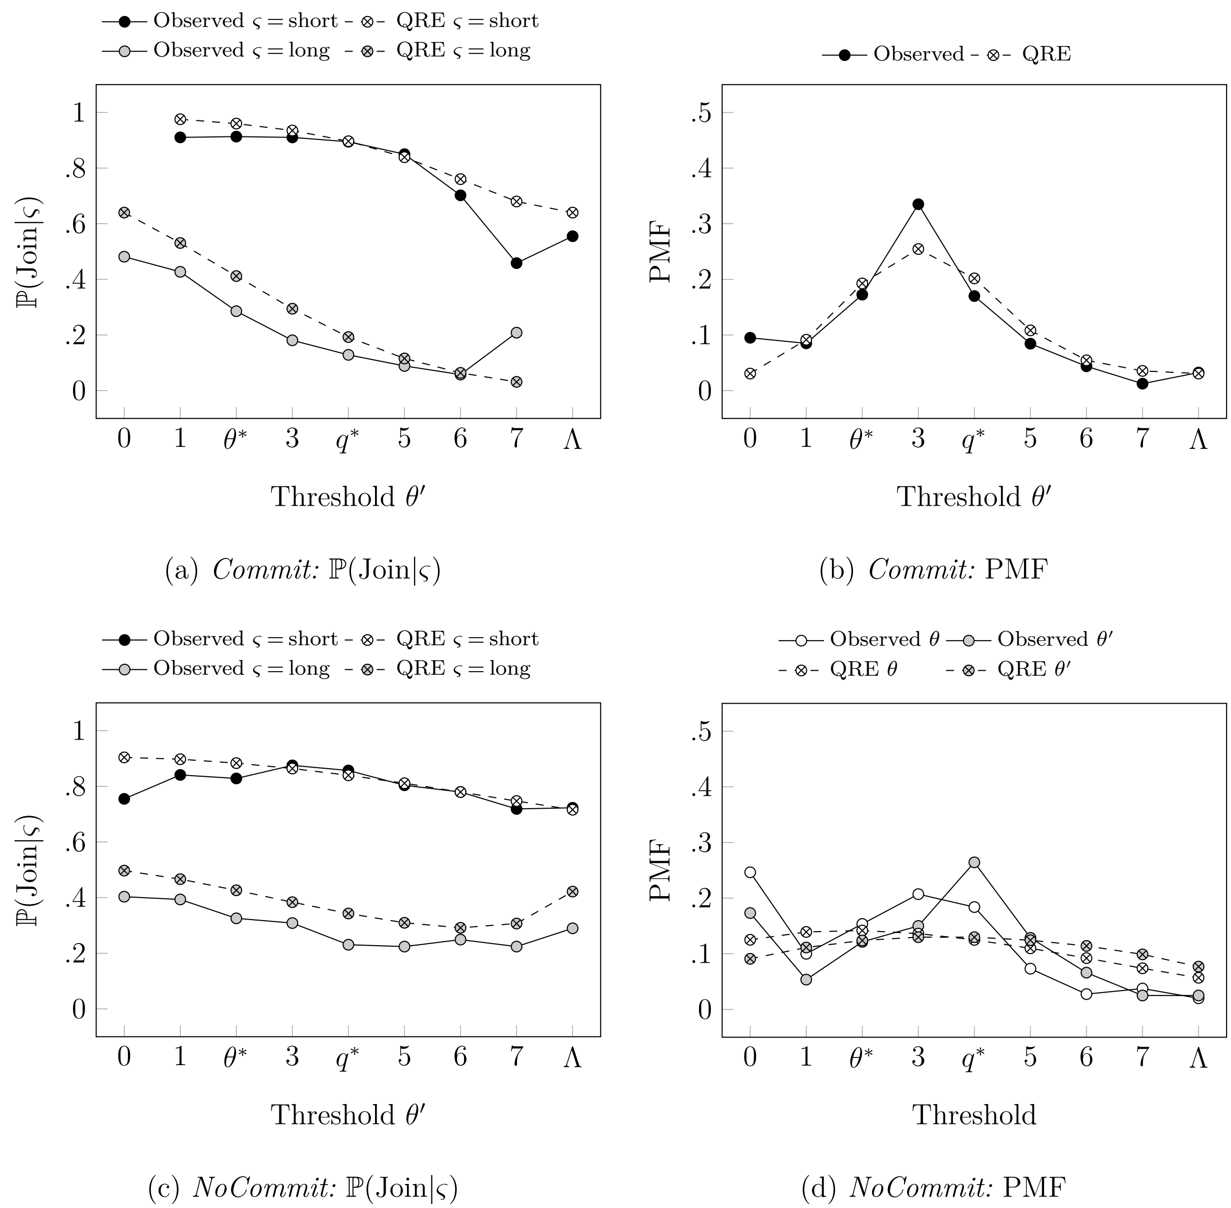

centralized admission control and price discrimination often are not feasible in social services settings, information design is a promising managerial lever to improve welfare. Leveraging their informational advantage about specific aspects of their (often complex) operational systems, service providers can mitigate the inefficient use of their resources by carefully and strategically communicating information in a way that persuades self-interested and rational low-need customers to forgo service for the benefit of others. Can efficient information design arise, in equilibrium, between providers and customers who may fall short of the rationality assumptions that commonly underlie theoretical analyses? This is an empirical question without an obvious answer, and one that is not easy to address because of the relatively high data demands that a rigorous test of theoretical predictions poses on the empirical setting. We present the results from laboratory experiments designed to shed first light on some of the key predictions and assumptions that information design theory makes.

Our main result, and key managerial implication for social service providers, is that careful information design improves social welfare. Importantly, our data shows that the service provider's ability to influence customer behaviour via credible signals does not hinge on their ability to formally commit to the information policy that drives such signals. This result is managerially relevant because commitment devices might be prohibitively costly to design in many real settings.

Our study is a first step towards empirically testing the efficacy of information design as a managerial lever to improve welfare in congestion-prone social service systems. But our design choices are not without limitations, and they point towards important work required to further test the robustness of our results and answer additional questions of practical relevance - can communication be effective if the provider does not share with its customers explicitly the policy that drives the signals; if the provider does share with its customers full information about the relevant system state, rather than curating the state into vague binary signals; if the provider can communicate to its customers an information policy that differs from the implemented one, without the risk of such deceit being detected? Opportunities for empirical research on the important topic of information design in social service operations abound.

References

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

Akşin, Z., B. Ata, S. M. Emadi, and C.-L. Su (2017). Impact of delay announcements in call centers: An empirical approach. *Operations Research 65*(1), 242–265.

Allon, G. and A. Bassamboo (2011). Buying from the babbling retailer? the impact of availability information on customer behavior. *Management Science 57*(4), 713–726.

Allon, G., A. Bassamboo, and I. Gurvich (2011). "we will be right with you": Managing customer expectations with vague promises and cheap talk. *Operations research 59*(6), 1382–1394.

Allon, G., A. Bassamboo, and R. S. Randhawa (2012). Price as a signal of product availability: Is it cheap? *Available at SSRN 3393502*.

Allon, G. and M. Kremer (2019). Behavioral foundations of queuing systems. In E. K. Karen Donohue and S. Leider (Eds.), *The Handbook of Behavioral Operations*, Handbooks in Operations Research and Management Science, pp. 325–366. Hoboken, NJ: John Wiley Sons, Inc.

Anand, K. S. and M. Goyal (2009). Strategic information management under leakage in a supply chain. *Management Science 55*(3), 438–452.

Ansari, S., L. Debo, M. Ibanez, S. Iravani, M. Malik, et al. (2022). Under-promising and over-delivering to improve patient satisfaction at emergency departments: Evidence from a field experiment providing wait information. *Available at SSRN 4135705*.

Antonides, G., V. P. and van Aalst (2002). Consumer perception and evaluation of waiting time: a field experiment. *Journal of Consumer Psychology*, 193—-202.

Anunrojwong, J., K. Iyer, and V. Manshadi (2022). Information design for congested social services optimal need-based persuasion. *Management Science*.

Aydinliyim, T., M. S. Pangburn, and E. Rabinovich (2017). Inventory disclosure in online retailing. *European Journal of Operational Research 261*(1), 195–204.

Beer, R., H.-S. Ahn, and S. Leider (2022). The impact of decision rights on innovation sharing. *Management Science 68*(11), 7793–8514.

Bergemann, D. and S. Morris (2019). Information design: A unified perspective. *Journal of Economic Literature 57*(1), 44–95.

Brandts, J., C. G. (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics 14*, 375–398.

Buell, R., E. Porter, and M. Norton (2021). Surfacing the submerged state: Operational transparency increases trust in and engagement with government. *Manufacturing & Service Operations Management*, 781–802.

Buell, R. W., T. Kim, and C.-J. Tsay (2017). Creating reciprocal value through operational transparency. *Management Science*, 1673–1695.

Buell, R. W. and M. I. Norton (2011). The labor illusion: How operational transparency increases perceived value. *Management Science*, 1564—-1579.

Cachon, G. P. and M. A. Lariviere (2001). Contracting to assure supply: How to share demand forecasts in a supply chain. *Management science 47*(5), 629–646.

Che, Y.-K. and O. Tercieux (2021). Optimal queue design. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 312–313.

Chen, D. L., M. Schonger, and C. Wickens (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

Chen, H. and M. Frank (2004). Monopoly pricing when customers queue. *IIE Transactions 36*(6), 569–581.

Chen, Y. and J. J. Hasenbein (2020). Knowledge, congestion, and economics: Paramater uncertainty in naor's model. *Queueuing Systems 96*, 83–99.

Chen, Y., X. Su, and X. Zhao (2012). Modeling bounded rationality in capacity allocation games with the quantal response equilibrium. *Management Science 58*(10), 1952–1962.

Chicco, D., N. Tötsch, and G. Jurman (2021). The matthews correlation coefficient (mcc) is more reliable than balanced accuracy, bookmaker informedness, and markedness in two-class confusion matrix evaluation. *BioData mining 14*(1), 1–22.

Cui, R. and H. Shin (2018). Sharing aggregate inventory information with customers: Strategic cross-selling and shortage reduction. *Management Science 64*(1), 381–400.

Debo, L. and G. Van Ryzin (2009). Creating sales with stock-outs. *Available at SSRN 1923706*.

Dong, J., E. Yom-Tov, and G. B. Yom-Tov (2019). The impact of delay announcements on hospital network coordination and waiting times. *Management Science 65*(5), 1969–1994.

Drakopoulos, K., S. Jain, and R. Randhawa (2021). Persuading customers to buy early: The value of personalized information provisioning. *Management Science 67*(2), 828–853.

Economou, A. (2021). The impact of information structure on strategic behavior in queueing systems. In *(2021) Queueing Theory 2, Advanced Trends.* Wiley/ISTE.

Edelson, N. M. and D. K. Hilderbrand (1975). Congestion tolls for poisson queuing processes. *Econometrica: Journal of the Econometric Society*, 81–92.

Fréchette, G. R., A. Lizzeri, and J. Perego (2022). Rules and commitment in communication: An experimental analysis. *Econometrica 90*(5), 2283–2318.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2016). *Quantal response equilibrium: A stochastic theory of games.* Princeton University Press.

Goldschmidt, K., M. Kremer, D. J. Thomas, and C. W. Craighead (2021). Strategic sourcing under severe disruption risk: Learning failures and under-diversification bias. *Manufacturing & Service Operations Management 23*(4), 761–780.

Hassin, R. (1986). Consumer information in markets with random product quality: The case of queues and balking. *Econometrica: Journal of the Econometric Society*, 1185–1195.

Haviv, M. and B. Oz (2018). Self-regulation of an unobservable queue. *Management Science 64*(5), 2380–2389.

Huang, T., G. Allon, and A. Bassamboo (2013). Bounded rationality in service systems. *Manufacturing & Service Operations Management 15*(2), 263–279.

Hui, M. K. and D. K. Tse (1996). What to tell consumers in waits of different lengths: An integrative model of service evaluation. *Journal of Marketing 60*(2), 81–90.

Ibrahim, R. (2018). Sharing delay information in service systems: a literature survey. *Queueing Systems 89*(1), 49–79.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Kremer, M. and F. de Vericourt (2023). (mis)managing diagnostic accuracy under congestion. *Operations Research*, 895–916.

Kremer, M. and L. Debo (2016). Inferring quality from wait time. *Management Science 62*(10), 3023–3038.

Lingenbrink, D. and K. Iyer (2019). Optimal signaling mechanisms in unobservable queues. *Operations research 67*(5), 1397–1416.

McFadden, D. (1981). Econometric models of probabilistic choice. *Structural analysis of discrete data with econometric applications 198272*.

Munichor, N. and A. Rafaeli (2007). Numbers or apologies? customer reactions to telephone waiting time fillers. *Journal of Applied Psychology 92*(2), 511.

Naor, P. (1969). The regulation of queue size by levying tolls. *Econometrica: journal of the Econometric Society*, 15–24.

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.

Schmidt, W., V. Gaur, R. Lai, and A. Raman (2015). Signaling to partially informed investors in the newsvendor model. *Production and Operations Management 24*(3), 383–401.

Shone, R., V. A. Knight, and J. E. Williams (2013). Comparisons between observable and unobservable m/m/1 queues with respect to optimal customer behavior. *European Journal of Operational Research 227*(1), 133–141.

Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management 10*(4), 566–589.

Webb, E., Q. Yu, and K. Bretthauer (2020). Linking delay announcements, abandonment, and service time. *Working Paper*.

Yu, M., H.-S. Ahn, and R. Kapuscinski (2015). Rationing capacity in advance selling to signal quality. *Management Science 61*(3), 560–577.

Yu, Q., G. Allon, and A. Bassamboo (2017). How do delay announcements shape customer behavior? an empirical study. *Management Science 63*(1), 1–20.

Yu, Q., Y. Zhang, and Y.-P. Zhou (2022). Delay information in virtual queues: A large-scale field experiment on a major ride-sharing platform. *Management Science 68*(8), 5745–5757.

Appendix

A. Technical Proofs

A.1. Proposition 1

PROOF. Recall that $\alpha_w \in [0,1]$ denotes a fixed joining probability for low-need customers. Based on this, we can compute the expected social welfare as follows:

$$\Omega = \mathbb{E}\left[\sum_{k=1}^{\Lambda} U_k\right] = r\mathbb{E}[J] - c\mathbb{E}\left[\sum_{k=1}^{J} k\right] = r\mathbb{E}[J] - \frac{c}{2}\left(\mathbb{E}[J^2] + \mathbb{E}[J]\right), \tag{6}$$

where J is a random variable that represents the total number of customers that *join* the system. Since all H type customers join, based on Equation (6) it follows that $J \sim \mathcal{B}(\Lambda, p_h + (1-p_h)\alpha_w)$ is a binomial random variable, such that $\mathbb{E}[J] = \Lambda(p_h + (1-p_h)\alpha_w)$, and $\mathbb{E}[J^2] = \Lambda(p_h + (1-p_h)\alpha_w) + \Lambda(\Lambda-1)(p_h + (1-p_h)\alpha_w)^2$. Based on this, after some algebra, we have that:

$$\frac{\partial \Omega}{\partial \alpha_w} = \Omega'(\alpha_w) = \Lambda(1-p_h)(r - c(1 + p_h(\Lambda-1)(1-\alpha_w) + \alpha_w(\Lambda-1))),$$

$$\Omega'(0) = \Lambda(1-p_h)(r - c(1 + p_h(\Lambda-1))),$$

$$\Omega'(1) = \Lambda(1-p_h)(r - c\Lambda) < 0,$$

$$\frac{\partial^2 \Omega}{\partial p_s^2} = \Omega''(\alpha_w) = -c(\Lambda-1)\Lambda(1-p_h)^2 < 0.$$

Since $\Omega'(\alpha_w)$ decreases strictly in α_w, and since $\Omega'(1) < 0$, we note that whenever $\Omega'(0) \leq 0 \iff \frac{r}{c} \leq 1 + p_h(\Lambda-1)$, the expected welfare in the system decreases for all values of α_w. This means that for $\frac{r}{c} \leq 1 + p_h(\Lambda-1)$, $\alpha_w^* = 0$ maximizes social welfare. Now, for the case $\frac{r}{c} > 1 + p_h(\Lambda-1)$, we have that $\Omega'(0) > 0$. Since $\Omega'(\alpha_w)$ decreases strictly in α_w, and since $\Omega'(1) < 0$, it follows that there is a unique $\alpha_w^* \in (0,1)$ that maximizes social welfare. In particular, such α_w^* satisfies $\Omega'(\alpha_w^*) = 0$. It is easy to see that in this case we have that $\alpha_w^* = \frac{r - c(1+p_h(\Lambda-1))}{c(\Lambda-1)(1-p_h)}$. ■

A.2. Proposition 2

PROOF. Consider a tagged L customer that joins with probability α', when the other L customers join with probability α. Since all H type customers always join, the tagged customer knows that a typical customer will join the system with probability $p_h + (1-p_h)\alpha$. When the tagged customer arrives, she does not know the number of people in the system, however, she can compute the expected number in equilibrium:

$$\mathbb{E}[Q] = \sum_{w=0}^{\Lambda-1} \mathbb{E}[Q|w]\mathbb{P}(w) = \frac{1}{\Lambda}\sum_{w=0}^{\Lambda-1} \mathbb{E}[Q|w],$$

where w is the remaining customers to arrive behind. Since customers are randomly ordered, they have w remaining customers behind them with probability $\frac{1}{\Lambda}$ for all $w \in \{0, 1, \cdots, \Lambda - 1\}$. Conditional on w, customers can arrive to a system with up to $\Lambda - 1 - w$ customers. For example, the first customer (i.e., $w = \Lambda - 1$) arrives to a system with 0 customers, and the last customer (i.e., $w = 0$) arrives to a system with 0 or 1 or 2, and so on up to $\Lambda - 1$ customers. Moreover, notice that, conditional on w, the probability to find q customers depends on how many customers joined in the previous times $\Lambda - 1, \Lambda - 2, \cdots, w + 1$; that is, in the previous $\Lambda - 1 - w$ times. Based on the above, we have that $Q|w \sim \mathcal{B}(\Lambda - 1 - w, p_h + (1 - p_h)\alpha)$ is a binomial random variable with expected value $\mathbb{E}[Q|w] = (\Lambda - 1 - w)(p_h + (1 - p_h)\alpha)$. With this, the tagged customer can compute the expected number of people in the system upon arrival:

$$\begin{aligned}
\mathbb{E}[Q] &= \frac{1}{\Lambda}\sum_{w=0}^{\Lambda-1}\mathbb{E}[Q|w] = \frac{1}{\Lambda}\sum_{w=0}^{\Lambda-1}(\Lambda - 1 - w)(p_h + (1 - p_h)\alpha) \\
&= \frac{p_h + (1 - p_h)\alpha}{\Lambda}\sum_{k=1}^{\Lambda-1} k = \frac{(p_h + (1 - p_h)\alpha)(\Lambda - 1)\Lambda}{2\Lambda} \\
&= \frac{(p_h + (1 - p_h)\alpha)(\Lambda - 1)}{2}.
\end{aligned}$$

Then, the tagged customer's expected utility is $\mathbb{E}[U(\alpha', \alpha)] = (1 - \alpha')0 + \alpha'(r - c(\mathbb{E}[Q] + 1))$. To find his best response against α, the tagged customer has to find the α' that maximizes $\mathbb{E}[U(\alpha', \alpha)]$. Note that the function $\mathbb{E}[U(\alpha', \alpha)]$ is linear with respect to α', so the tagged customer bases her decision on the sign of the quantity $r - c(\mathbb{E}[Q] + 1)$. Let the root of $r - c(\mathbb{E}[Q] + 1) = 0$ be

$$\bar{\alpha} = \frac{2r - c(2 + p_h(\Lambda - 1))}{c(\Lambda - 1)(1 - p_h)}.$$

Based on the above, the set of best responses against α, $BR(\alpha) = \arg\max_{\alpha'}\mathbb{E}[U(\alpha', \alpha)]$, is given by

$$BR(\alpha) = \begin{cases} 0 & \text{if } \alpha > \bar{\alpha} \\ [0, 1] & \text{if } \alpha = \bar{\alpha} \\ 1 & \text{if } \alpha < \bar{\alpha} \end{cases}$$

We can now proceed to the computation of the equilibrium strategies:

- The strategy $(\alpha^* = 0)$ is an equilibrium strategy, if and only if $0 \in BR(0)$, i.e., $0 \geq \bar{\alpha}$, which reduces to $\frac{(\Lambda-1)p_h}{2} + 1 \geq \frac{r}{c}$.
- The strategy $(\alpha^* = 1)$ is an equilibrium strategy if, and only if, $1 \in BR(1)$, i.e., $1 \leq \bar{\alpha}$, which reduces to $\frac{\Lambda-1}{2} + 1 \leq \frac{r}{c}$.

- The strategy $\alpha^* \in (0,1)$ is an equilibrium strategy if, and only if, $\alpha^* \in BR(\alpha^*)$, i.e., $\alpha^* = \bar{\alpha}$. This is valid so long as $\bar{\alpha} \in (0,1)$ which occurs if and only if $\frac{(\Lambda-1)p_h}{2} + 1 < \frac{r}{c} < \frac{(\Lambda-1)}{2} + 1$.

■

A.3. Corollary 1

Proof. From Propositions 1 and 2, we first note that since $\frac{(\Lambda-1)p_h}{2} + 1 < 1 + p_h(\Lambda - 1)$, if $\alpha^* = 0$ then $\alpha_w^* = 0$.

From Proposition 1, we consider the joining probability $\alpha_w^* = \frac{r - c(1+p_h(\Lambda-1))}{c(\Lambda-1)(1-p_h)} < 1$, and from Proposition 2 we consider the mixing probability $\alpha^* = \frac{2r - c(2+p_h(\Lambda-1))}{c(\Lambda-1)(1-p_h)}$. After some simple algebra, it is easy to see that $\alpha^* > \alpha_w^* \iff r > c$, which is our assumption by construction. This also means that for $\alpha^* = 1$ we have that $\alpha^* > \alpha_w^*$. ■

A.4. Lemma 1.

Proof. We analyze the case for customers and the service provider separately.

Customers' Best Response. Let $\theta \in \{0, \cdots, \Lambda\}$ be a fixed implemented threshold. Consider a tagged customer that joins with probabilities α_s' and α_l', when the rest of customers join with probabilities α_s and α_l. Given the parameter space in the game, $\frac{r}{c} \geq \frac{\Lambda-1}{2} + 1$, it follows that $r - c(\mathbb{E}[Q] + 1) \geq 0$ for any $\alpha_s \in [0,1]$ and $\alpha_l \in [0,1]$. Indeed, for the case in which all customers join $\alpha_s = \alpha_l = 1$, we have that $\mathbb{E}[Q] = \frac{\Lambda-1}{2}$. Since $\mathbb{E}[Q|\varsigma = s] \leq \mathbb{E}[Q]$ for any threshold θ, it follows that $r - c(\mathbb{E}[Q|\varsigma = s] + 1) \geq 0$ for any $\alpha_s \in [0,1]$, $\alpha_l \in [0,1]$, and threshold θ. This means that it is always in the best interest of the tagged customer to join the system with probability 1 upon receiving a *s*hort signal, irrespective of how others behave. It follows that $\alpha_s^* = 1$.

Now, when customers receive a *l*ong signal, they know that they are in a system with $q \in \{\theta, \theta+1, \cdots, \Lambda-1\}$ customers. This is because L type customers that receive a *s*hort signal join the system with probability $\alpha_s^* = 1$ in equilibrium. Based on this, and since H types always join, the tagged customer can condition on the indexes and compute the expected number of customers that arrived after the threshold θ. For example, conditional on having an index that coincides with the threshold θ, the tagged customer knows that she is in a system with $q = \theta$ customers with probability 1. Conditional on having an index that coincides with $\theta + 1$, the tagged customer knows that she is in a system with $q = \{\theta, \theta + 1\}$ customers. The probability of being in either of these states depends on the number customers N that arrived after the threshold and before the tagged customer. For this case, we have that $\mathbb{P}(Q = \theta) = \mathbb{P}(N = 0)$, and $\mathbb{P}(Q = \theta + 1) = \mathbb{P}(N = 1)$, where

$N \sim \mathcal{B}(1, p_h + (1-p_h)\alpha_l)$ is a binomial random variable. Thus, conditional on having an index that coincides with $\theta+1$ a customer expects to find $\theta + \mathbb{E}[N] = \theta + p_h + (1-p_h)\alpha_l$ customers.

More generally, conditional on having an index that coincides with $\theta + j$ $(j \in \{0, 1, \cdots, \Lambda - \theta - 1\})$, the tagged customer knows that she in a system with $q = \{\theta, \theta+1, \cdots, \theta+j\}$ customers, with respective probabilities $\mathbb{P}(N=0), \mathbb{P}(N=1), \cdots, \mathbb{P}(N=j)$, with $N \sim \mathcal{B}(j, p_h + (1-p_h)\alpha_l)$. And thus, conditional on having an index that coincides with $\theta + j$, a customer expects to find $\theta + \mathbb{E}[N] = \theta + j(p_h + (1-p_h)\alpha_l)$ customers. Based on this, the tagged customer computes:

$$
\begin{aligned}
\mathbb{E}[Q|\varsigma = l] &= \frac{1}{\Lambda - \theta}(\theta + (\theta + (p_h + (1-p_h)\alpha^l)) + (\theta + 2(p_h + (1-p_h)\alpha_l)) \\
&+ \cdots + (\theta + (\Lambda - \theta - 1)(p_h + (1-p_h)\alpha_l))) \\
&= \frac{1}{\Lambda - \theta}(\theta(\Lambda - \theta) + \frac{(p_h + (1-p_h)\alpha_l)(\Lambda - \theta - 1)(\Lambda - \theta)}{2}) \\
&= \frac{2\theta + (p_h + (1-p_h)\alpha_l)(\Lambda - \theta - 1)}{2}.
\end{aligned}
$$

Then, the tagged customer's expected utility is $\mathbb{E}[U(\alpha_l', \alpha_l)] = (1-\alpha_l')0 + \alpha_l'(r - c(\mathbb{E}[Q|\varsigma = l] + 1))$. To find her best response against α_l, the tagged customer has to find the α_l' that maximizes $\mathbb{E}[U(\alpha_l', \alpha_l)]$. Note that the function $\mathbb{E}[U(\alpha_l', \alpha_l)]$ is linear with respect to α_l', so the tagged customer bases her decision on the sign of the quantity $r - c(\mathbb{E}[Q|\varsigma = l] + 1)$. Let the root of $r - c(\mathbb{E}[Q|\varsigma = l] + 1) = 0$ be

$$
\bar{\alpha_l} = \frac{2r - c(2(qr+1) + p_h(\Lambda - \theta - 1))}{c(1-p_h)(\Lambda - \theta - 1)}.
$$

The set of best responses of the tagged customer against α_l, $BR(\alpha_l) = \arg\max_{\alpha_l'} \mathbb{E}[U(\alpha_l', \alpha_l)]$, is given by

$$
BR(\alpha_l) = \begin{cases} 0 & \text{if } \alpha_l > \bar{\alpha}_l \\ [0,1] & \text{if } \alpha_l = \bar{\alpha}_l \\ 1 & \text{if } \alpha_l < \bar{\alpha}_l \end{cases}
$$

We can now proceed to the computation of the customer equilibrium strategies, which represent the best response of all customers to a given θ:

- The strategy $(\alpha_l^* = 0)$ upon receiving a *l*ong signal is an equilibrium strategy if, and only if, $0 \in BR(0)$, i.e., $0 \geq \bar{\alpha}_l$, which reduces to $\theta \geq \frac{2r - 2c - cp_h(\Lambda - 1)}{c(2-p_h)}$. Recall that $\bar{\theta} = \lceil \frac{2r - 2c - cp_h(\Lambda - 1)}{c(2-p_h)} \rceil$. That is, we have $\alpha_l^* = 0$ whenever $\theta \in \{\bar{\theta}, \cdots, \Lambda - 1\}$.

- The strategy ($\alpha_l^* = 1$) upon receiving a *l*ong signal is an equilibrium strategy if, and only if, $1 \in BR(1)$, i.e., $1 \leq \bar{\alpha}_l$, which reduces to $\theta \leq \frac{2r-c(\Lambda+1)}{C}$. Let $\bar{\bar{\theta}} = \lfloor \frac{2r-c(\Lambda+1)}{c} \rfloor$, we have that $\alpha_l^* = 1$ whenever $\theta \in \{0, \cdots, \bar{\bar{\theta}}\}$.
- The strategy $\alpha_l^* \in (0,1)$ is an equilibrium strategy if, and only if, $\alpha_l^* \in BR(\alpha_l^*)$, i.e., $\alpha_l^* = \bar{\alpha}_l$. This is valid so long as $\bar{\alpha}_l \in (0,1)$ which occurs if, and only if, $\frac{2r-c(\Lambda+1)}{c} < \theta < \frac{2r-2c-cp_h(\Lambda-1)}{c(2-p_h)}$, that is, $\bar{\bar{\theta}} < \theta < \bar{\theta}$.

Provider's Best Response. From the previous analysis, we see that $\alpha_s = 1$ irrespective of θ. We can thus restrict attention to that case in what follows. Let $\alpha_s = 1, \alpha_l \in [0,1]$ be a fixed customer joining behaviour. Since in this case all customers join when receiving a *s*hort signal, we let $J(\theta) = \theta + N$ be a random variable that represents the total number of customers that *join* the system under a given threshold θ, where $N \sim \mathcal{B}(\Lambda - \theta, p_h + (1-p_h)\alpha_l)$ is a binomial random variable that represents the total number of customers that join after the threshold θ. Based on this, the *expected social welfare*, plugging in $\alpha_s = 1$, is:

$$\begin{aligned}\Omega(\theta,\alpha_l) &= r\mathbb{E}[J(\theta)] - \frac{c}{2}(\mathbb{E}[J(\theta)^2] + \mathbb{E}[J(\theta)]) \\ &= r(\theta + \mathbb{E}[N]) - \frac{c}{2}(\theta^2 + 2\theta\mathbb{E}[N] + \mathbb{E}[N^2] + \theta + \mathbb{E}[N]),\end{aligned}$$

where $\mathbb{E}[N] = (\Lambda-\theta)(p_h + (1-p_h)\alpha_l)$, and $\mathbb{E}[N^2] = (\Lambda-\theta)(p_h + (1-p_h)\alpha_l) + (\Lambda-\theta)(\Lambda-\theta-1)(p_h + (1-p_h)\alpha_l)^2$. Now, consider the following expression:

$$\begin{aligned}&\Omega(\theta,\alpha_l) - \Omega(\theta-1,\alpha_l) = \\ &\quad (1-p_h-(1-p_h)\alpha_l)(r - c\Lambda(p_h+(1-p_h)\alpha_l) - c(1-p_h-(1-p_h)\alpha_l)\theta). \qquad (7)\end{aligned}$$

First, from (7), we note that if $\alpha_l = 1$ then we have that $\Omega(\theta,1) = \Omega(\theta-1,1)$ for any threshold θ. That is, any threshold $\theta^* \in \{0, \cdots, \Lambda\}$ represents the provider's best response if $\alpha_s = 1, \alpha_l = 1$. This is intuitive since in this case all customers join the system with probability 1 irrespective of the received signal, such that changing the threshold (while holding customer behaviour constant) does not have an impact on social welfare. Now, from (7), we note that if $p_h \geq \frac{r}{c\Lambda}$ (which is the case by construction), for any $\alpha_l \in [0,1)$, we have that $\Omega(\theta,\alpha_l) - \Omega(\theta-1,\alpha_l) \leq 0$. This implies that for sufficiently large p_h, for a given customer response $\alpha_l \in [0,1)$, the social welfare is at its maximum at a threshold $\theta = 0$. That is, a threshold $\theta^* = 0$ represents the provider's best response if $\alpha_s = 1, \alpha_l < 1$. ■

A.5. Proposition 3.

Proof. From Lemma 1, we know that customers best respond with $\alpha_s^* = 1, \alpha_l^* = 1$ to any threshold $\theta \in \{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$. Also, we know that if customers join with $\alpha_s = 1, \alpha_l = 1$, then the provider best responds with any threshold $\theta^* \in \{0, \cdots, \Lambda\}$. This implies that $\theta^* \in \{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$ and $\alpha_s^*(\theta^*) = 1, \alpha_l^*(\theta^*) = 1$ constitute equilibria in our game. Finally, from Lemma 1, we know that customers best respond with $\alpha_s^* = 1, \alpha_l^* < 1$ to any threshold $\theta \in \{\bar{\bar{\theta}} + 1, \cdots, \Lambda - 1\}$. Also, we know that if customers join with $\alpha_s = 1, \alpha_l < 1$, then the provider best responds with a threshold $\theta^* = 0$. This implies that none of the thresholds $\theta \in \{\bar{\bar{\theta}}+1, \cdots, \Lambda-1\}$, nor a customer behaviour of the form $\alpha_s = 1, \alpha_l < 1$ can arise in equilibrium. ■

A.6. Proposition 4.

Proof. We analyze the case with and without commitment separately.

With commitment. Based on Lemma 1, we consider three different cases separately: (1) $\theta \le \bar{\bar{\theta}}$, (2) $\theta \ge \bar{\theta}$, and (3) $\bar{\bar{\theta}} < \theta < \bar{\theta}$.

Case: $\theta \le \bar{\bar{\theta}}$. From Lemma 1, a threshold $\theta \le \bar{\bar{\theta}}$ induces a customer equilibrium where every L type customer decides to join irrespective of the signal, i.e., $\alpha_s^* = 1, \alpha_l^* = 1$. We will see that a threshold in $\theta \ge \bar{\theta}$ allows to improve social welfare in comparison to this case where all customers join, and thus $\theta \le \bar{\bar{\theta}}$ does not arise in equilibrium.

Case: $\theta \ge \bar{\theta}$. From Lemma 1, a threshold $\theta \ge \bar{\theta}$ induces a customer equilibrium where L type customers that receive a *s*hort signal join with probability $\alpha_s^* = 1$, and those that receive a *l*ong signal join with probability $\alpha_l^* = 0$. To compute the social welfare in the system, let $J(\theta) = \theta + N$ be a random variable that represents the total number of customers that *join* the system in equilibrium (recall all customers join when receiving a *s*hort signal since $\alpha_s^* = 1$), where $N \sim \mathcal{B}(\Lambda - \theta, p_h + (1 - p_h)\alpha_l^*)$ is a binomial random variable that represents the total number of customers that join after a given threshold θ. Based on this, the *expected social welfare* is:

$$\begin{aligned}\Omega(\theta) &= r\mathbb{E}[J(\theta)] - \frac{c}{2}(\mathbb{E}[J(\theta)^2] + \mathbb{E}[J(\theta)]) \\ &= r(\theta + \mathbb{E}[N]) - \frac{c}{2}(\theta^2 + 2\theta\mathbb{E}[N] + \mathbb{E}[N^2] + \theta + \mathbb{E}[N]),\end{aligned}$$

where $\mathbb{E}[N] = (\Lambda - \theta)p_h$, and $\mathbb{E}[N^2] = (\Lambda - \theta)p_h + (\Lambda - \theta)(\Lambda - \theta - 1)p_h^2$. Now, consider the following expression:

$$\Omega(\theta) - \Omega(\theta - 1) = (1 - p_h)(r - c\Lambda p_h - c(1 - p_h)\theta).$$

Since we have by construction that $p_h \geq \frac{r}{c\Lambda}$, it follows that $\Omega(\theta)-\Omega(\theta-1)<0$ for all thresholds θ. It follows that the expected social welfare decreases strictly for all $\theta \geq \bar{\theta}$. This implies that from the thresholds such that $\theta \geq \bar{\theta}$, the threshold $\bar{\theta}$ yields the highest social welfare. Moreover, we note that with a threshold $\theta=\Lambda$, all customers join the system (as in the previous described case for $\theta \leq \bar{\bar{\theta}}$). Since the expected social welfare decreases strictly for all $\theta \geq \bar{\theta}$, it follows that the social welfare achieved at $\bar{\theta}$ is higher than that achieved with any threshold in $\theta \leq \bar{\bar{\theta}}$.

Case: $\bar{\bar{\theta}} < \theta < \bar{\theta}$. From Lemma 1, a threshold $\bar{\bar{\theta}} < \theta < \bar{\theta}$ induces a customer equilibrium where L type customers that receive a *s*hort signal join $\alpha_s^*=1$, and those that receive a *l*ong signal join with probability $\alpha_l^*=\frac{2r-c(2(\theta+1)+p_h(\Lambda-\theta-1))}{c(1-p_h)(\Lambda-\theta-1)}$. To compute the social welfare in the system, let $J(\theta)=\theta+N$ be a random variable that represents the total number of customers that *join* the system in equilibrium (recall all customers join when receiving a *s*hort signal since $\alpha_s^*=1$), where $N\sim\mathcal{B}(\Lambda-\theta, p_h+(1-p_h)\alpha_l^*)$ is a binomial random variable that represents the total number of customers that join after a given threshold θ. Based on this, the *expected social welfare* in the system can be written as:

$$\begin{aligned}\Omega(\theta) &= r\mathbb{E}[J(\theta)] - \frac{c}{2}(\mathbb{E}[J(\theta)^2]+\mathbb{E}[J(\theta)]) \\ &= r(\theta+\mathbb{E}[N]) - \frac{c}{2}(\theta^2+2\theta\mathbb{E}[N]+\mathbb{E}[N^2]+\theta+\mathbb{E}[N]),\end{aligned}$$

where $\mathbb{E}[N]=(\Lambda-\theta)(p_h+(1-p_h)\alpha_l^*)$, and $\mathbb{E}[N^2]=(\Lambda-\theta)(p_h+(1-p_h)\alpha_l^*)+(\Lambda-\theta)(\Lambda-\theta-1)(p_h+(1-p_h)\alpha_l^*)^2$. Now, consider the following expression:

$$\Omega(\theta)-\Omega(\theta-1)=r-c\theta.$$

Since it is easy to see that $\bar{\theta}\leq q^*=\lfloor\frac{r}{c}\rfloor$, it follows that $\Omega(\theta)-\Omega(\theta-1)>0$. This means that the expected social welfare increases in θ whenever $\bar{\bar{\theta}}<\theta<\bar{\theta}$, and thus from all the threshold such that $\bar{\bar{\theta}}<\theta<\bar{\theta}$, the threshold $\theta=\bar{\theta}-1$ achieves the maximum social welfare.

Finally, based on the above, it remains to see which of the two thresholds, $\bar{\theta}-1$ or $\bar{\theta}$, achieves a higher social welfare. For this, consider the following inequality:

$$\begin{aligned}&\Omega(\bar{\theta})-\Omega(\bar{\theta}-1)\geq 0\\ \iff &2(r-c\bar{\theta})+(\Lambda-\bar{\theta})p_h(2r-c(2(1+\bar{\theta})+p_h(\Lambda-\bar{\theta}-1)))\geq 0.\end{aligned}$$

We let $g(\bar{\theta})\doteq 2(r-c\bar{\theta})+(\Lambda-\bar{\theta})p_h(2r-c(2(1+\bar{\theta})+p_h(\Lambda-\bar{\theta}-1)))$. We note that the quantity $r-c\bar{\theta}\geq 0 \iff \bar{\theta}\leq\frac{r}{c}$ since $\bar{\theta}\leq q^*$, and that the quantity $2r-c(2(1+\bar{\theta})+p_h(\Lambda-\bar{\theta}-1))\leq 0 \iff \bar{\theta}\geq\frac{2r-2c-cp_h(\Lambda-1)}{c(2-p_h)}$, since

$\bar{\theta} = \lceil \frac{2r-2c-cp_h(\Lambda-1)}{c(2-p_h)} \rceil$. It follows that depending on the parameters in the system it is possible for $g(\bar{\theta}) < 0$, $g(\bar{\theta}) = 0$, or $g(\bar{\theta}) > 0$. It follows that if $g(\bar{\theta}) \geq 0$, we have that $\theta^* = \bar{\theta}$ and $\alpha_l^*(\bar{\theta}) = 0$ arise in equilibrium. Otherwise, we have that $\theta^* = \bar{\theta} - 1$ and $\alpha_l^*(\bar{\theta} - 1) \in (0, 1)$.

Without commitment. We can show that communicated thresholds θ' cannot affect the equilibrium outcomes of Proposition 3. Customers would take θ' into consideration only if such communicated thresholds are credible. That is, if the service provider does not have the incentive to implement a threshold θ that differs from θ' given customers' response to it ($\alpha_s^*(\theta') = 1, \alpha_l^*(\theta') \in [0, 1]$) as described in Lemma 1. Formally, for θ' to be credible, it is required that the expected social welfare, for $\alpha_s^*(\theta') = 1, \alpha_l^*(\theta') \in [0, 1]$, complies with $\Omega(\theta') \geq \Omega(\theta)$ for all possible thresholds θ. Consider the following expression:

$$\begin{aligned} &\Omega(\theta') - \Omega(\theta' - 1) = \\ &\quad (1 - p_h - (1 - p_h)\alpha_l^*(\theta'))(r - c\Lambda(p_h + (1 - p_h)\alpha_l^*(\theta')) - c(1 - p_h - (1 - p_h)\alpha_l^*(\theta'))\theta'). \end{aligned} \tag{8}$$

First, from (8), we note that if $\theta' \in \{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$, then we have that $\Omega(\theta') = \Omega(\theta' - 1)$ for any threshold θ'. This implies that $\theta' \in \{0, \cdots, \bar{\bar{\theta}}, \Lambda\}$ are credible thresholds. However, since $\alpha_s^*(\theta') = 1, \alpha_l^*(\theta') = 1$, the actual values of the thresholds θ' do not affect the equilibrium outcomes of Proposition 3. Finally, from (8), we note that if $\theta' \in \{\bar{\bar{\theta}} + 1, \cdots, \Lambda - 1\}$, then we have that $\Omega(\theta') < \Omega(\theta' - 1)$ for any threshold θ'. This implies that $\theta' \in \{\bar{\bar{\theta}} + 1, \cdots, \Lambda - 1\}$ are not credible since the service provider has the incentive to implement $\theta = 0$ in this case. The fact that communicated thresholds do not affect customer behaviour is intuitive since customers can infer the equilibrium θ with strategic thinking, and since the communication of θ' is costless to the service provider. Based on this, since customers are not influenced by communicated thresholds, the provider is indifferent in the selection of θ' and, thus, selects them randomly. This resulting uninformative babble regarding the implemented threshold, and customer disregard of the communicated threshold represent mutual best responses between the customers and the service provider. ■

A.7. Corollary 2

Proof. We analyze the case with and without commitment separately.

With commitment. In this case, from Proposition 4 we know that L type customers join with probability 1 when receiving a *s*hort signal, and with probability 0 when receiving a *l*ong signal in equilibrium. Moreover, we know that the service provider selects a threshold $0 < \theta^* < q^*$ in equilibrium, and that $\alpha_s^*(\theta^*) = 1$ and $\alpha_l^*(\theta^*) \in [0, 1)$.

Based on this, it is easy to see that $\mathbb{P}(\text{Join}|Q<q^*)=\frac{\alpha_s^*(\theta^*)\mathbb{P}(Q<\theta^*)+\alpha_l^*(\theta^*)\mathbb{P}(Q\in\{\theta^*,\cdots,q^*-1\})}{\mathbb{P}(Q<q^*)}\in(0,1)$. Also, we can see that $\mathbb{P}(\text{Balk}|Q\geq q^*)=1-\mathbb{P}(\text{Join}|Q\geq q^*)=1-\frac{\alpha_l^*(\theta^*)\mathbb{P}(Q\geq q^*)}{\mathbb{P}(Q\geq q^*)}=1-\alpha_l^*(\theta^*)\in(0,1]$. Finally, from this it follows that $\mathcal{I}=\mathbb{P}(\text{Join}|Q<q^*)+\mathbb{P}(\text{Balk}|Q\geq q^*)-1=\frac{\mathbb{P}(\text{Join}|Q<\theta^*)(\alpha_s^*(\theta^*)-\alpha_l^*(\theta^*))}{\mathbb{P}(\text{Join}|Q<q^*)}\in(0,1)$.

Without commitment. In this case, from Proposition 4 we know that L type customers join with probability 1 in equilibrium, irrespective of the true state q. Based on this, it is easy to see that $\mathbb{P}(\text{Join}|Q<q^*)=\mathbb{P}(\text{Join})=1,\ \mathbb{P}(\text{Balk}|Q\geq q^*)=\mathbb{P}(\text{Balk})=0$. Based on this and expression (1), it follows that $\mathcal{I}=\mathbb{P}(\text{Join}|Q<q^*)+\mathbb{P}(\text{Balk}|Q\geq q^*)-1=0$. ■

B. How to Compute Expected/Implied Metrics

We focus on the performance optimality of participant choices. To reduce the effect of noise (i.e., variability due to the random matching of customer choices and realizations of customer types and order of arrivals) we compute the expected values of diverse metrics given the participant strategies in our data. Given the design of our experiment, it is natural to compute these implied metrics at the round-cohort level. We will see that in the *Commit* and *NoCommit* treatments, the computation of most of these metrics requires the input of the implemented threshold θ_{jtc}. This input is, by definition, an integer number and, thus, averaging across rounds or cohorts may generate average thresholds that are not integer-valued. We now describe how to compute the expected values of these metrics for all the treatments.

B.1. NoSignal treatment

In this treatment, in each round $t \in \{1,\dots,T\}$, for each cohort $c \in \{1,\dots,C\}$, and for each customer $i \in \{1,\cdots,\Lambda\}$, we observe their strategy $A_{itc} \equiv \big\{a_{itc} \in \{1,0\}\big\}$, where 1 represents the *Join* action and 0 the *Balk* action. Based on this, we compute $\varphi_{tc} = \frac{1}{\Lambda}\sum_i a_{itc}$, that is, the implied joining proportion in a particular round and cohort. Note that φ_{tc} is not necessarily the same as the *realized* joining proportion in our data, due to the random matching of customer choices and realizations of customer types.

Social Welfare. Let $J_{tc} \sim \mathcal{B}(\Lambda, p_h + (1-p_h)\varphi_{tc})$ be a binomial random variable that represents the number of customers that join in a given round and cohort. It follows that the expected social welfare is given by:

$$\Omega_{tc} = \mathbb{E}[rJ_{tc} - c\sum_{k=1}^{J_{tc}} k] = r\mathbb{E}[J_{tc}] - \frac{c}{2}\mathbb{E}[J_{tc}(J_{tc}+1)] = r\mathbb{E}[J_{tc}] - \frac{c}{2}(\mathbb{E}[J_{tc}^2] + \mathbb{E}[J_{tc}]),$$

with $\mathbb{E}[J_{tc}] = \Lambda(p_h + (1-p_h)\varphi_{tc})$, and $\mathbb{E}[J_{tc}^2] = \Lambda(p_h + (1-p_h)\varphi_{tc}) + \Lambda(\Lambda-1)(p_h + (1-p_h)\varphi_{tc})^2$.

Overall Joining Probability of L Types. In this case, it simply follows that $\mathbb{P}_{tc}(Join) = \varphi_{tc}$.

Joining Probability of L Types given that $q < q^*$. Recall that q^* represents the Naor threshold. In this treatment, it simply follows that $\mathbb{P}_{tc}(Join|q < q^*) = \varphi_{tc}$. Indeed, since customers do not have any information about the queue state when they make their decisions, their joining is independent of queue states.

Balking Probability of L Types given that $q \geq q^*$. In this treatment, it simply follows that $\mathbb{P}_{tc}(Balk|q \geq q^*) = 1 - \varphi_{tc}$. Similarly, since customers do not have any information about the queue state when they make their decisions, their balking is independent of queue states.

Queue-length Informativeness. We recall that $\mathcal{I}_{tc} = \mathbb{P}_{tc}(Join|q < q^*) + \mathbb{P}_{tc}(Balk|q \geq q^*) - 1$. In this treatment, it is easy to see that $\mathcal{I}_{tc} = 0$.

B.2. Commit and NoCommit treatment

In these treatments, in each round $t \in \{1, \ldots, T\}$, for each cohort $c \in \{1, \ldots, C\}$, we observe the provider j's implemented threshold decision θ_{jtc} (*Commit*,*NoCommit*), and the communicated threshold decision θ'_{jtc} (*NoCommit*). For each customer $i \in \{1, \cdots, \Lambda\}$, we observe their strategy $A_{itc} \equiv \{a_{itc}(\varsigma = s|\theta_{jtc}), a_{itc}(\varsigma = l|\theta_{jtc})\}$ (*Commit*), and $A_{itc} \equiv \{a_{itc}(\varsigma = s|\theta'_{jtc}), a_{itc}(\varsigma = l|\theta'_{jtc})\}$ (*NoCommit*), where $a_{itc} \in \{0, 1\}$, such that 1 represents the *Join* action and 0 the *Balk* action. Based on this, in the *Commit* treatment, we compute $\varphi^s_{tc} = \frac{1}{\Lambda}\sum_i a_{itc}(\varsigma = s|\theta_{jtc})$ and $\varphi^l_{tc} = \frac{1}{\Lambda}\sum_i a_{itc}(\varsigma = l|\theta_{jtc})$, that is, the implied joining proportion for a given message in a particular round and cohort. Similarly, in the *NoCommit* treatment, we compute $\varphi^s_{tc} = \frac{1}{\Lambda}\sum_i a_{itc}(\varsigma = s|\theta'_{jtc})$ and $\varphi^l_{tc} = \frac{1}{\Lambda}\sum_i a_{itc}(\varsigma = l|\theta'_{jtc})$. We note that φ^s_{tc} and φ^l_{tc} are not necessarily the same as the *realized* joining proportions in our data, due to the random matching of customer choices and realizations of customer types and order of arrivals.

System Dynamics. In both *Commit* and *NoCommit* treatment, for a given set of φ^s_{tc}, φ^l_{tc} and θ_{jtc}, we can capture the expected dynamics in the queue with the state probabilities $\mathbb{P}_{tc}(q, w)$. We have that $q \in \{0, 1, \cdots, \Lambda - 1\}$ represents the number of customers in queue that an arriving tagged customer encounters, and $w \in \{\Lambda - 1, \cdots, 1, 0\}$ represents the remaining number of customers to arrive after an arriving tagged customer. We let $\sigma_{tc}(q, w)$ be the probability $\mathbb{P}_{tc}(\varsigma = s|q, w)$, that is, the probability that a L signal is sent in state (q, w). Also, we have that $1 - \sigma_{tc}(q, w)$ is the probability $\mathbb{P}_{tc}(\varsigma = l|q, w)$, that is, the probability that a H signal is sent in state (q, w).

Notice that state (q, w) can be reached only from states $(q-1, w+1)$ and $(q, w+1)$. If the system is in state $(q-1, w+1)$, then the system transitions into state (q, w) with probability $p_h + (1-p_h)(\sigma(q-1, w+1)\varphi^s_{tc} + (1-\sigma_{tc}(q-1, w+1))\varphi^l_{tc})$. On the other hand, if the system is in state $(q, w+1)$, then the system transitions into state (q, w) with probability $(1-p_h)(\sigma_{tc}(q-1, w+1)(1-\varphi^s_{tc}) + (1-\sigma_{tc}(q-1, w+1))(1-\varphi^l_{tc}))$. Based on the above, we can capture the queueing dynamics with the following recursive expression:

$$\mathbb{P}_{tc}(q, w) = \mathbb{P}_{tc}(q-1, w+1)(p_h + (1-p_h)(\sigma_{tc}(q-1, w+1)\varphi^s_{tc} + (1-\sigma_{tc}(q-1, w+1))\varphi^l_{tc}))$$

$$+\mathbb{P}_{tc}(q,w+1)(1-p_h)(\sigma_{tc}(q-1,w+1)(1-\varphi^s_{tc})$$
$$+(1-\sigma_{tc}(q-1,w+1))(1-\varphi^l_{tc})),$$

with boundary conditions $\mathbb{P}_{tc}(0,\Lambda-1)=1/\Lambda$, $\mathbb{P}_{tc}(0,w)=\mathbb{P}_{tc}(0,w+1)(1-p_h)(\sigma_{tc}(0,w+1)(1-\varphi^s_{tc})+(1-\sigma_{tc}(0,w+1))(1-\varphi^l_{tc}))$ for all $w<\Lambda-1$, and $\mathbb{P}_{tc}(q,w)=\mathbb{P}_{tc}(q-1,w+1)(p_h+(1-p_h)(\sigma_{tc}(q-1,w+1)\varphi^s_{tc}+(1-\sigma_{tc}(q-1,w+1))\varphi^l_{tc}))$ for all $q>0,w<\Lambda-1$ such that $q+w=\Lambda-1$.

Note that, given the fixed threshold structure of the signalling mechanism, we have that $\sigma_{tc}(q,w)=\sigma_{tc}(q)$ for all w, and that $\sigma_{tc}(q)=1$ if $q<\theta_{jtc}$ and $\sigma_{tc}(q)=0$ otherwise. Finally, we note that given the set of φ^s_{tc}, φ^l_{tc}, and θ_{jtc}, in the *NoCommit* treatment, the system dynamics do not depend on the announced θ'_{jtc}.

Social Welfare. We consider the value function $V_{tc}(q,w)$. For a given state (q,w), the expected future utility $V_{tc}(q,w)$ is equal to the immediate expected utility, $p_h(r-c(q+1))+(1-p_h)(r-c(q+1))(\sigma_{tc}(q,w)\varphi^s_{tc}+(1-\sigma_{tc}(q,w))\varphi^l_{tc})$, plus the expected utility from time $w-1$ onward, $(p_h+(1-p_h)(\sigma_{tc}(q,w)\varphi^s_{tc}+(1-\sigma_{tc}(q,w))\varphi^l_{tc}))V_{tc}(q+1,w-1)+(1-p_h)((\sigma_{tc}(q,w)(1-\varphi^s_{tc})+(1-\sigma_{tc}(q,w))(1-\varphi^l_{tc})))V(q,w-1)$. After some simple algebra we can simplify the expression of the expected future utility:

$$V_{tc}(q,w)=(r-c(q+1)+V_{tc}(q+1,w-1))(p_h+(1-p_h)(\sigma_{tc}(q,w)\varphi^s_{tc}$$
$$+(1-\sigma_{tc}(q,w))\varphi^l_{tc}))+(1-p_h)((\sigma_{tc}(q,w)(1-\varphi^s_{tc})$$
$$+(1-\sigma_{tc}(q,w))(1-\varphi^l_{tc})))V_{tc}(q,w-1),$$

with boundary conditions $V_{tc}(q,0)=(r-c(q+1))(p_h+(1-p_h)(\sigma_{tc}(q,w)\varphi^s_{tc}+(1-\sigma_{tc}(q,w))\varphi^l_{tc}))$ for all q. Notice that, based on the recursive nature of the value function $V_{tc}(q,w)$, the expected social welfare in the system, Ω_{tc}, is given by $V_{tc}(0,\Lambda-1)$. Note that, given the fixed threshold structure of the signalling mechanism, we have that $\sigma_{tc}(q,w)=\sigma_{tc}(q)$ for all w, and that $\sigma_{tc}(q)=1$ if $q<\theta_{jtc}$ and $\sigma_{tc}(q)=0$ otherwise.

Joining Probability of L Types for a given message. In this case, we simply have that $\mathbb{P}_{tc}(Join|\varsigma=l)=\varphi^l_{tc}$ and that $\mathbb{P}_{tc}(Join|\varsigma=s)=\varphi^s_{tc}$.

Overall Joining Probability of L Types. In this case, we can condition on the received message as follows:

$$\mathbb{P}_{tc}(Join)=\mathbb{P}_{tc}(Join|\varsigma=s)\mathbb{P}_{tc}(\varsigma=s)+\mathbb{P}_{tc}(Join|\varsigma=l)\mathbb{P}_{tc}(\varsigma=l)$$
$$=\varphi^s_{tc}\mathbb{P}_{tc}(\varsigma=s)+\varphi^l_{tc}(1-\mathbb{P}_{tc}(\varsigma=s)).$$

Finally, we can compute the probabilities for a message based on the state probabilities as $\mathbb{P}_{tc}(\varsigma = s) = \sum_w \sum_q \sigma_{tc}(q,w)\mathbb{P}_{tc}(q,w)$, and $\mathbb{P}_{tc}(\varsigma = l) = 1 - \mathbb{P}_{tc}(\varsigma = s)$.

Joining Probability of L Types given that $q < q^*$. In this case, we can condition on the received message as follows:

$$\begin{aligned}
\mathbb{P}_{tc}(Join|q<q^*) &= \mathbb{P}_{tc}(Join|q<q^*,\varsigma=s)\mathbb{P}_{tc}(\varsigma=s|q<q^*) + \mathbb{P}_{tc}(Join|q<q^*,\varsigma=l)\mathbb{P}_{tc}(\varsigma=l|q<q^*) \\
&= \mathbb{P}_{tc}(Join|\varsigma=s)\mathbb{P}_{tc}(\varsigma=s|q<q^*) + \mathbb{P}_{tc}(Join|\varsigma=l)\mathbb{P}_{tc}(\varsigma=l|q<q^*) \\
&= \varphi_{tc}^s\mathbb{P}_{tc}(\varsigma=s|q<q^*) + \varphi_{tc}^l\mathbb{P}_{tc}(\varsigma=l|q<q^*) \\
&= \varphi_{tc}^s\sum_q \mathbb{P}_{tc}(\varsigma=s|q<q^*,q)\mathbb{P}_{tc}(q|q<q^*) + \varphi_{tc}^l\sum_q \mathbb{P}_{tc}(\varsigma=l|q<q^*,q)\mathbb{P}_{tc}(q|q<q^*) \\
&= \varphi_{tc}^s\sum_q \mathbb{P}_{tc}(\varsigma=s|q)\mathbb{P}_{tc}(q|q<q^*) + \varphi_{tc}^l\sum_q \mathbb{P}_{tc}(\varsigma=l|q)\mathbb{P}_{tc}(q|q<q^*) \\
&= \varphi_{tc}^s\sum_{q=0}^{q^*-1}\sigma_{tc}(q)\frac{\mathbb{P}_{tc}(q)}{\sum_{q=0}^{q^*-1}\mathbb{P}_{tc}(q)} + \varphi_{tc}^l\sum_{q=0}^{q^*-1}(1-\sigma_{tc}(q))\frac{\mathbb{P}_{tc}(q)}{\sum_{q=0}^{q^*-1}\mathbb{P}_{tc}(q)} \\
&= \frac{\sum_{q=0}^{q^*-1}(\varphi_{tc}^s\sigma_{tc}(q) + \varphi_{tc}^l(1-\sigma_{tc}(q)))\mathbb{P}_{tc}(q)}{\sum_{q=0}^{q^*-1}\mathbb{P}_{tc}(q)} \\
&= \frac{\sum_{q=0}^{q^*-1}(\varphi_{tc}^s\sigma_{tc}(q) + \varphi_{tc}^l(1-\sigma_{tc}(q)))\sum_w\mathbb{P}_{tc}(q,w)}{\sum_{q=0}^{q^*-1}\sum_w\mathbb{P}_{tc}(q,w)}.
\end{aligned}$$

We recall that $\sigma_{tc}(q) = 1$ if $q < \theta_{jtc}$ and $\sigma_{tc}(q) = 0$ otherwise.

Balking Probability of L Types given that $q \geq q^*$. In this case, we can condition on the received message as follows:

$$\begin{aligned}
\mathbb{P}_{tc}(Balk|q\geq q^*) &= \mathbb{P}_{tc}(Balk|q\geq q^*,\varsigma=s)\mathbb{P}_{tc}(\varsigma=s|q\geq q^*) + \mathbb{P}_{tc}(Balk|q\geq q^*,\varsigma=l)\mathbb{P}_{tc}(\varsigma=l|q\geq q^*) \\
&= \mathbb{P}_{tc}(Balk|\varsigma=s)\mathbb{P}_{tc}(\varsigma=s|q\geq q^*) + \mathbb{P}_{tc}(Balk|\varsigma=l)\mathbb{P}_{tc}(\varsigma=l|q\geq q^*) \\
&= (1-\varphi_{tc}^s)\mathbb{P}_{tc}(\varsigma=s|q\geq q^*) + (1-\varphi_{tc}^l)\mathbb{P}_{tc}(\varsigma=l|q\geq q^*) \\
&= (1-\varphi_{tc}^s)\sum_q \mathbb{P}_{tc}(\varsigma=s|q\geq q^*,q)\mathbb{P}_{tc}(q|q\geq q^*) + (1-\varphi_{tc}^l)\sum_q \mathbb{P}_{tc}(\varsigma=l|q\geq q^*,q)\mathbb{P}_{tc}(q|q\geq q^*) \\
&= (1-\varphi_{tc}^s)\sum_q \mathbb{P}_{tc}(\varsigma=s|q)\mathbb{P}_{tc}(q|q\geq q^*) + (1-\varphi_{tc}^l)\sum_q \mathbb{P}_{tc}(\varsigma=l|q)\mathbb{P}_{tc}(q|q\geq q^*) \\
&= (1-\varphi_{tc}^s)\sum_{q=q^*}^{\Lambda-1}\sigma_{tc}(q)\frac{\mathbb{P}_{tc}(q)}{\sum_{q=q^*}^{\Lambda-1}\mathbb{P}_{tc}(q)} + (1-\varphi_{tc}^l)\sum_{q=q^*}^{\Lambda-1}(1-\sigma_{tc}(q))\frac{\mathbb{P}_{tc}(q)}{\sum_{q=q^*}^{\Lambda-1}\mathbb{P}_{tc}(q)} \\
&= \frac{\sum_{q=q^*}^{\Lambda-1}((1-\varphi_{tc}^s)\sigma_{tc}(q) + (1-\varphi_{tc}^l)(1-\sigma_{tc}(q)))\mathbb{P}_{tc}(q)}{\sum_{q=q^*}^{\Lambda-1}\mathbb{P}_{tc}(q)}
\end{aligned}$$

$$= \frac{\sum_{q=q^*}^{\Lambda-1}((1-\varphi_{tc}^s)\sigma_{tc}(q)+(1-\varphi_{tc}^l)(1-\sigma_{tc}(q)))\sum_w \mathbb{P}_{tc}(q,w)}{\sum_{q=q^*}^{\Lambda-1}\sum_w \mathbb{P}_{tc}(q,w)}.$$

We recall that $\sigma_{tc}(q)=1$ if $q<\theta_{jtc}$ and $\sigma_{tc}(q)=0$ otherwise.

Queue-length Informativeness. This is simply computed as $\mathcal{I}_{tc}=\mathbb{P}_{tc}(Join|q<q^*)+\mathbb{P}_{tc}(Balk|q\geq q^*)-1$.

Persuasiveness of Signals. We recall that $\mathcal{P}_{ct}=\mathbb{P}_{tc}(Balk|\varsigma=l)$ if $\theta_{jtc}<q^*$. It simply follows that:

$$\mathcal{P}_{ct}=\begin{cases}1-\varphi_{tc}^l & \text{if } \theta_{jtc}<q^*,\\ \text{not defined} & otherwise.\end{cases}$$

Note that a natural metric to compare against in the *NoSignals* treatment (as a baseline) is $\mathbb{P}_{tc}(Balk|q<q^*)=1-\mathbb{P}_{tc}(Join|q<q^*)=1-\varphi_{tc}$.

C. Experiment

Table 4 displays the main realized and implied results for total welfare, service provider behaviour, and customer behaviour.

Table 4 Experimental Results (standard deviations in parentheses)

Treatment	Realized Metrics						Implied Metrics					
	Welfare	$\mathbb{P}$(Join)	$\mathbb{P}(\text{Join} \mid q < q^*)$	$\mathbb{P}(\text{Balk} \mid q \geq q^*)$	$\mathcal{I}$	$\mathcal{P}$	Welfare	$\mathbb{P}$(Join)	$\mathbb{P}(\text{Join} \mid q < q^*)$	$\mathbb{P}(\text{Balk} \mid q \geq q^*)$	$\mathcal{I}$	$\mathcal{P}$
NoSignals	172.33	0.55	0.53	0.42	-0.05	-	172.84	0.57	0.57	0.43	0	-
	(11.34)	(0.05)	(0.06)	(0.06)	(0.05)	-	(22.59)	(0.09)	(0.09)	(0.09)	(0.00)	-
Commit	199.38	0.45	0.63	0.77	0.41	0.74	200.70	0.48	0.66	0.74	0.41	0.74
	(8.03)	(0.03)	(0.07)	(0.06)	(0.10)	(0.07)	(10.86)	(0.05)	(0.06)	(0.09)	(0.12)	(0.09)
NoCommit	194.61	0.48	0.58	0.68	0.26	0.64	198.70	0.48	0.61	0.68	0.29	0.66
	(14.56)	(0.07)	(0.14)	(0.09)	(0.19)	(0.08)	(10.20)	(0.05)	(0.11)	(0.09)	(0.18)	(0.09)

C.1. Social Welfare

Table 5 presents OLS regressions regarding social welfare.

Table 5 OLS Regressions: Social Welfare

	Realized Social Welfare				Implied Social Welfare			
	(1a)	(2a)	(3a)†	(4a)	(1b)	(2b)	(3b)†	(4b)
(Intercept)	146.46***	146.46***	179.23***	146.46***	157.80***	157.80***	188.40***	157.80***
	(9.56)	(9.43)	(10.65)	(9.26)	(14.57)	(14.36)	(6.97)	(14.10)
Commit	35.60*	-	2.83	35.60**	33.81$^\cdot$	-	3.20	33.81*
	(11.59)	-	(12.42)	(11.22)	(14.96)	-	(7.72)	(14.48)
NoCommit	-	32.77*	-	32.77*	-	30.61$^\cdot$	-	30.61$^\cdot$
	-	(14.30)	-	(14.05)	-	(15.99)	-	(15.71)
Round	1.26**	1.26**	0.75*	1.26**	0.73**	0.73**	0.50*	0.73***
	(0.35)	(0.35)	(0.31)	(0.34)	(0.16)	(0.15)	(0.20)	(0.15)
Commit∗Round	-0.42	-	0.09	-0.42	-0.29	-	-0.06	-0.29
	(0.54)	-	(0.50)	(0.52)	(0.31)	-	(0.33)	(0.30)
NoCommit∗Round	-	-0.51	-	-0.51	-	-0.23	-	-0.23
	-	(0.46)	-	(0.46)	-	(0.26)	-	(0.25)
N	640	800	960	1200	640	800	960	1200
R^2	0.04	0.02	0.01	0.02	0.16	0.11	0.03	0.11

$^\cdot p < 0.1$, $^* p < 0.05$, $^{**} p < 0.01$, $^{***} p < 0.001$.

Each data point consists of the total social welfare achieved by a cohort in a given round and session. Thus $N =$ Number of sessions∗2(cohorts)∗40(rounds).

Standard errors clustered at the session level.

† : Models (3a) and (3b) do not consider the *NoSignal* treatment, and take as baseline the *NoCommit* treatment to compare it with the *Commit* treatment.

C.2. Customers' Choices

Table 6 presents logistic regressions regarding customers' strategies.

Table 6 Logistic Regressions: Customer Strategies

	$\mathbb{P}$(Join|Short Wait)			$\mathbb{P}$(Join|Long Wait)		
	(1a)	(2a)	(3a)	(1b)	(2b)	(3b)
(Intercept)	3.61***	1.55**	2.80***	0.11	-0.09	0.05
	(0.72)	(0.58)	(0.54)	(0.33)	(0.39)	(0.40)
NoCommit	-	-	-0.82*	-	-	-0.11
	-	-	(0.35)	-	-	(0.28)
Communicated Threshold	-0.14*	0.03	-0.14*	-0.38***	-0.18***	-0.37***
	(0.07)	(0.05)	(0.06)	(0.06)	(0.04)	(0.06)
*NoCommit * Communicated Threshold*	-	-	0.18*	-	-	0.20**
	-	-	(0.08)	-	-	(0.07)
Round	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
Gender.M	-0.82**	-0.13	-0.38	-0.14	-0.16	-0.15
	(0.28)	(0.35)	(0.26)	(0.16)	(0.30)	(0.20)
N	3200	4480	7680	3200	4480	7680

$^{\cdot}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Each data point represents the answer (Join/Not Join) by a Customer in a given round.

Models (1a) and (1b) are restricted to the *Commit* treatment, and models (2a) and (2b) are restricted to the *NoCommit* treatment.

Standard errors clustered at the session level.

C.3. Service providers' Choices

Table 7 presents OLS regressions regarding service providers' lying behaviour.

C.4. Post-Experimental Survey Data

At the end of the experiment, all participants respond to a survey relevant to their experiences in the experiment. Participants answered with a number between 1 to 7 (where 1 represents *strongly disagree*, 3 *neutral*, and 7 *strongly agree*) to the following statements:

For customers:

- Q1: I was in good mood
- Q2: I felt in control of my outcomes
- Q3: The service provider and the customers shared the same objective
- Q4: I understood the service provider's information strategy (*Commit*, *NoCommit*)

Table 7 OLS Regressions: Service Providers' Lying Behaviour

	Realized Social Welfare			Implied Social Welfare		
	(1a)	(2a)	(3a)	(1b)	(2b)	(3b)
(Intercept)	168.75***	188.75***	174.42***	186.54***	199.31***	189.47***
	(13.68)	(17.17)	(15.91)	(8.46)	(11.10)	(9.18)
Lie	10.75***	-	-	8.40***	-	-
	(1.12)	-	-	(0.79)	-	-
$\|Lie\|$	-	0.74	-	-	2.40$^{\cdot}$	-
	-	(2.12)	-	-	(1.13)	-
$Lie.Type(\theta' > \theta)$	-	-	24.42*	-	-	23.97**
	-	-	(7.94)	-	-	(5.12)
$Lie.Type(\theta' < \theta)$	-	-	-53.57***	-	-	-27.45**
	-	-	(7.74)	-	-	(4.62)
Round	0.43	0.74$^{\cdot}$	0.41	0.25	0.48$^{\cdot}$	0.27
	(0.33)	(0.33)	(0.29)	(0.25)	(0.24)	(0.22)
Gender.M	5.99	-6.85	7.42	0.65	-8.92	-0.36
	(5.01)	(6.10)	(6.56)	(2.91)	(4.63)	(3.54)
N	560	560	560	560	560	560
R^2	0.07	0.01	0.07	0.32	0.06	0.26

$^{\cdot}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Each data point consists of the total social welfare achieved by a service provider in a given round in the *NoComit* treatment.

Standard errors clustered at the session level.

The variable *Lie* is defined as $Lie = \theta' - \theta$, and $Lie.Type(\theta' > \theta)$, $Lie.Type(\theta' < \theta)$ represent indicator variables.

In models (3a) and (3b) the baseline for $Lie.Type$ is $Lie.Type(\theta' = \theta)$, such that the coefficients for $Lie.Type(\theta' > \theta)$ and $Lie.Type(\theta' < \theta)$ correspond to differences when service providers lie in comparison to when they are honest.

- Q5: The service provider's information strategy had an impact on my decisions (*Commit*, *NoCommit*)
- Q6: I trusted the information strategy (*Commit*, *NoCommit*)

For service providers:

- Q1: I was in good mood
- Q2: I felt in control of my outcomes
- Q3: The service provider and the customers shared the same objective

- Q4: I understood the service provider's information strategy (*Commit*, *NoCommit*)
- Q5: My information strategy had an impact on customers' decisions (*Commit*, *NoCommit*)
- Q6: The information strategy was trusted by customers (*Commit*, *NoCommit*)

Table 8 presents the average results at the treatment level for each of the described statements above.

Table 8 Average Survey Responses at Treatment Level

Statement	Customers			Service Providers		
	NoSignals	*Commit*	*NoCommit*	*NoSignals*	*Commit*	*NoCommit*
Q1	4.69	4.45	4.54	4.83	4.2	4.07
Q2	3.10	3.38	3.22	2.33	3.6	3.21
Q3	3.21	3.55	3.43	3.33	3.1	3.86
Q4	-	4.39	3.94	-	4.1	5
Q5	-	5.65	4.41	-	4.6	4.64
Q6	-	3.98	3.28	-	2.6	4.14

Notes: Results are rounded to 2 decimals places.

D. Quantal Response Equilibrium (QRE)

We let $\pi_i(a_i, \mathbf{a_{-i}})$ denote player i's payoff from taking action a_i over an action set S_i when the rest of players choose actions $\mathbf{a_{-i}}$ over a vector of corresponding action sets $\mathbf{S_{-i}}$. In the QRE framework, players choose a *noisy* best response by maximizing $\mathbb{E}[\pi_i(a_i, \mathbf{A_{-i}})] + \epsilon_i$ instead of $\mathbb{E}[\pi_i(a_i, \mathbf{A_{-i}})]$, where ϵ_i is a noise term and $\mathbf{A_{-i}}$ represents a vector of random variables. By assuming that the ϵ_i terms are distributed i.i.d according to a mean-zero Gumbel distribution with scale parameter, $\beta > 0$, and with CDF, $F_i(z_i) = e^{-e^{-\frac{1}{\beta_i} z_i - \gamma}}$, where $\gamma \approx 0.5772$ is the Euler-Mascheroni constant, we obtain the following logit specification for choice probabilities (McFadden 1981):

$$\mathbb{P}_i(a_i) = \frac{e^{\mathbb{E}[\pi_i(a_i, \mathbf{A_{-i}})]/\beta_i}}{\sum_{a_i' \in S_i} e^{\mathbb{E}[\pi_i(a_i', \mathbf{A_{-i}})]/\beta_i}}. \tag{9}$$

A QRE of the game represents a strategy profile where all players choose distributions over actions in a consistent way, such that players have correct expectations and beliefs about the probability distributions of others. In the above expression (9), β_i captures player i's level of *bounded rationality* (Chen et al. 2012, Huang et al. 2013). This is because the parameter β_i is proportional to the standard deviation of the noise term ϵ_i ($\approx 0.78\beta_i$): When $\beta_i \to 0$, player i chooses the payoff-maximizing alternative with certainty (i.e., the theoretical prediction under full rationality). At the other extreme, when $\beta_i \to \infty$, player i lacks the ability to make any rational judgement and thus randomizes over all alternatives with equal probabilities. In what follows, for sake of parsimony, we use a common parameter β for every customer, and β_m for the service provider. We note that since high-need customers experience disutility $-\infty$ when balking (i.e., they do not have an outside option), they always join also under QRE.

D.1. No Signals Treatment

In this case, L type customers join with probability:

$$\varphi(\beta) = \frac{e^{(r - c(\mathbb{E}[Q]+1))/\beta}}{1 + e^{(r - c(\mathbb{E}[Q]+1))/\beta}}. \tag{10}$$

In this expression, customers compute the expected numbed of customers in the system $\mathbb{E}[Q]$ upon their arrival, based on their beliefs about how others behave. In equilibrium, customers' expectations and beliefs are consistent with behaviour. That is, the QRE of the game represents the solution of the fixed point problem given by equation (10), denoted by $\varphi^*(\beta)$. To compute the QRE, we first need to understand how to compute customers' expectations.

How to compute $\mathbb{E}[Q]$**.** We can compute the expected number of customers that an arriving customer encounters as follows:

$$\mathbb{E}[Q] = \sum_{q=0}^{\Lambda-1} q\mathbb{P}(q) = \sum_{q=0}^{\Lambda-1} q \sum_{w=0}^{\Lambda-1} \mathbb{P}(q,w),$$

where q represents the current number of customers in queue, and w the remaining customers to arrive. To compute the state probabilities $\mathbb{P}(q,w)$, we notice that states (q,w) can be reached only from states $(q-1,w+1)$ and $(q,w+1)$. For a given $\varphi(\beta)$, if the system is in state $(q-1,w+1)$, then the system transitions into state (q,w) with probability $p_h + (1-p_h)\varphi(\beta)$. On the other hand, if the system is in state $(q,w+1)$, the system transitions into state (q,w) with probability $(1-p_h)(1-\varphi(\beta))$. Based on this, we can capture the queueing dynamics with the following recursive expression:

$$\mathbb{P}(q,w) = \mathbb{P}(q-1,w+1)(p_h + (1-p_h)\varphi(\beta)) + \mathbb{P}(q,w+1)(1-p_h)(1-\varphi(\beta)),$$

with boundary conditions $\mathbb{P}(0,\Lambda-1) = 1/\Lambda$, $\mathbb{P}(0,w) = \mathbb{P}(0,w+1)(1-p_h)(1-\varphi(\beta))$ for all $w < \Lambda-1$, and $\mathbb{P}(q,w) = \mathbb{P}(q-1,w+1)(p_h + (1-p_h)\varphi(\beta))$ for all $q > 0, w < \Lambda-1$ such that $q+w = \Lambda-1$.

Estimation of Relevant Parameter. Recall that the QRE of the game is given by $\varphi^*(\beta)$. Let f denote the data set in this treatment, where we recall that in each round $t \in \{1,\cdots,T\}$ and for each cohort $c \in \{1,\cdots,C\}$, such data set contains for each customer $i \in \{1,\cdots,\Lambda\}$, their strategy $A_{itc} \equiv \{a_{itc}\}$, where $a_{itc} \in \{0,1\}$ is such that 1 represents the *Join* action and 0 the *Balk* action. Based on this, the log-likehood function is given by:

$$\mathcal{L}(\beta|f) = \sum_i \sum_t \sum_c \left(a_{itc}\log(\varphi^*(\beta)) + (1-a_{itc})\log(1-\varphi^*(\beta))\right).$$

We compute the value $\hat{\beta}$ that maximizes the above function $\mathcal{L}(\beta|f)$. That is, we compute the maximum-likelihood estimate for the parameter β, based on our experimental data.

Computation of Implied Aggregate Metrics Under Estimated QRE. In Appendix B.1, we describe how to compute the implied values of different aggregate metrics (e.g., expected social welfare Ω, informativeness $\mathcal{I}$) for a given joining proportion φ_{tc}. Similarly, to compute the implied values of the aggregate metrics under the estimated QRE, we use the same procedure described in Appendix B.1, but we use instead the estimated QRE joining proportion $\varphi^*(\hat{\beta})$.

D.2. Commitment Treatment

In this case, given an implemented threshold θ and a signal ς, L type customers join with probability:

$$\varphi_s(\theta;\beta) = \frac{e^{(r-c(\mathbb{E}[Q|\varsigma=s,\theta]+1))/\beta}}{1+e^{(r-c(\mathbb{E}[Q|\varsigma=s,\theta]+1))/\beta}} \quad \text{for} \quad \theta = 1,\cdots,\Lambda, \tag{11}$$

$$\varphi_l(\theta;\beta) = \frac{e^{(r-c(\mathbb{E}[Q|\varsigma=l,\theta]+1))/\beta}}{1+e^{(r-c(\mathbb{E}[Q|\varsigma=l,\theta]+1))/\beta}} \quad \text{for} \quad \theta = 0,\cdots,\Lambda-1. \tag{12}$$

In these expressions, customers compute the expected numbed of customers in the system $\mathbb{E}[Q|\varsigma,\theta]$ upon their arrival, based on the provider's implemented threshold θ and their beliefs about how other customers behave. Note that such expectation is not a function of the provider's parameter β_m (as in the case of *NoCommit* below) since with commitment, we have a sequential game where customers react to a known realized implemented threshold θ. This is irrespective of the fact that the provider selects thresholds in a noisy fashion. Formally, the service provider selects a probability distribution $\vartheta(\theta;\beta,\beta_m)$ over the choice of the implemented threshold θ:

$$\vartheta(\theta;\beta,\beta_m) = \frac{e^{(\Omega(\varphi_s(\theta;\beta),\varphi_l(\theta;\beta),\theta))/\beta_m}}{\sum_\theta e^{(\Omega(\varphi_s(\theta;\beta),\varphi_l(\theta;\beta),\theta))\beta_m}} \quad \text{for} \quad \theta = 0,\cdots,\Lambda. \tag{13}$$

Based on the above, in equilibrium, customers' and the service provider's expectations and beliefs are consistent with customers' and the service provider's behaviour. That is, the QRE of the game is given by the solution of the system of equations (11) - (13), denoted by $\varphi_s^*(\theta;\beta)$ for $\theta \in \{1,\cdots,\Lambda\}$, $\varphi_l^*(\theta;\beta)$ for $\theta \in \{0,\cdots,\Lambda-1\}$ and $\vartheta^*(\theta;\beta,\beta_m)$ for $\theta \in \{0,\cdots,\Lambda\}$. In this case, given the sequential nature of the game that arises from commitment, it is easy to see that we can first compute the *customer equilibrium*, $\varphi_s^*(\theta;\beta)$ and $\varphi_l^*(\theta;\beta)$, for any given threshold θ. That is, the solution of the system of equations given by (11) and (12). Then, based on such customer equilibrium, we can compute $\vartheta^*(\theta;\beta,\beta_m)$ by plugging the customer equilibrium directly in (13). This follows since the social welfare $\Omega(\cdot)$ does not depend on the probability $\vartheta(\theta;\beta,\beta_m)$, such that $\vartheta(\theta;\beta,\beta_m)$ represents the equilibrium distribution $\vartheta^*(\theta;\beta,\beta_m)$. To compute the QRE, we first need to understand how to compute customers' expectations and the service provider's expected social welfare.

How to compute $\mathbb{E}[Q|\varsigma,\theta]$**.** For a fixed implemented threshold θ, we can compute the expected number of customers as follows:

$$\mathbb{E}[Q|\varsigma,\theta] = \sum_{q=0}^{\Lambda-1} q\mathbb{P}(q|\varsigma,\theta) = \sum_{q=0}^{\Lambda-1} q\frac{\mathbb{P}(\varsigma|q,\theta)\mathbb{P}(q|\theta)}{\mathbb{P}(\varsigma|\theta)}$$

$$
\begin{aligned}
&= \sum_{q=0}^{\Lambda-1} q \frac{(\sum_{w=0}^{\Lambda-1} \mathbb{P}(\varsigma|q,w,\theta)\mathbb{P}(w|q,\theta))\mathbb{P}(q|\theta)}{\mathbb{P}(\varsigma|\theta)} \\
&= \sum_{q=0}^{\Lambda-1} q \frac{\sum_{w=0}^{\Lambda-1} \mathbb{P}(\varsigma|q,w,\theta)\mathbb{P}(q,w|\theta)}{\sum_{w=0}^{\Lambda-1}\sum_{q=0}^{\Lambda-1} \mathbb{P}(\varsigma|q,w,\theta)\mathbb{P}(q,w|\theta)},
\end{aligned}
$$

where q represents the current number of customers in queue, and w the remaining customers to arrive. To compute the state probabilities $\mathbb{P}(q,w|\theta)$, we need to consider the dynamics in the queue. Let $\mathbb{P}(\varsigma = s|q,w,\theta) = \sigma_\theta(q,w)$, and $\mathbb{P}(\varsigma = l|q,w,\theta) = 1-\sigma_\theta(q,w)$. First, notice that states (q,w) can be reached only from states $(q-1,w+1)$ and $(q,w+1)$. For a given $\varphi_s(\theta)$, $\varphi_l(\theta)$ and θ, if the system is in state $(q-1,w+1)$, then the system transitions into state (q,w) with probability $p_h + (1-p_h)(\sigma_\theta(q-1,w+1)\varphi_s(\theta) + (1-\sigma_\theta(q-1,w+1))\varphi_l(\theta))$. On the other hand, if the system is in state $(q,w+1)$, then the system transitions into state (q,w) with probability $(1-p_h)(\sigma_\theta(q-1,w+1)(1-\varphi_s(\theta)) + (1-\sigma_\theta(q-1,w+1))(1-\varphi_l(\theta)))$. Based on this, we can capture the queueing dynamics with the following recursive expression:

$$
\begin{aligned}
\mathbb{P}(q,w|\theta) =& \mathbb{P}(q-1,w+1|\theta)(p_h + (1-p_h)(\sigma_\theta(q-1,w+1)\varphi_s(\theta) + (1-\sigma_\theta(q-1,w+1))\varphi_l(\theta))) \\
&+ \mathbb{P}(q,w+1|\theta)(1-p_h)(\sigma_\theta(q-1,w+1)(1-\varphi_s(\theta)) + (1-\sigma_\theta(q-1,w+1))(1-\varphi_l(\theta))),
\end{aligned}
$$

with boundary conditions $\mathbb{P}(0,\Lambda-1|\theta) = 1/\Lambda$, $\mathbb{P}(0,w|\theta) = \mathbb{P}(0,w+1|\theta)(1-p_h)(\sigma_\theta(0,w+1)(1-\varphi_s(\theta)) + (1-\sigma_\theta(0,w+1))(1-\varphi_l(\theta)))$ for all $w < \Lambda-1$, and $\mathbb{P}(q,w|\theta) = \mathbb{P}(q-1,w+1|\theta)(p_h + (1-p_h)(\sigma_\theta(q-1,w+1)\varphi_s(\theta) + (1-\sigma_\theta(q-1,w+1))\varphi_l(\theta)))$ for all $q>0, w<\Lambda-1$ such that $q+w=\Lambda-1$. Note that given the fixed threshold structure of the signalling mechanism we have that $\sigma_\theta(q,w) = \sigma_\theta(q)$ for all w, and that $\sigma_\theta(q) = 1$ if $q<\theta$ and $\sigma_\theta(q) = 0$ otherwise.

***How to compute* $\boldsymbol{\Omega}$.** To compute the service provider's expected social welfare, we consider the value function $V(q,w|\theta)$. For a given $\varphi_s(\theta)$, $\varphi_l(\theta)$ and θ, in state (q,w), the expected future utility $V(q,w|\theta)$ is equal to the immediate expected utility, $p_h(r-c(q+1)) + (1-p_h)(r-c(q+1))(\sigma_\theta(q,w)\varphi_s(\theta) + (1-\sigma_\theta(q,w))\varphi_l(\theta))$, plus the expected utility from time $w-1$ onward, $(p_h + (1-p_h)(\sigma_\theta(q,w)\varphi_s(\theta) + (1-\sigma_\theta(q,w))\varphi_l(\theta))V(q+1,w-1|\theta) + (1-p_h)((\sigma_\theta(q,w)(1-\varphi_s(\theta)) + (1-\sigma_\theta(q,w))(1-\varphi_l(\theta))))V(q,w-1|\theta)$. It follows that:

$$
\begin{aligned}
V(q,w|\theta) =& (r-c(q+1) + V(q+1,w-1|\theta))(p_h + (1-p_h)(\sigma_\theta(q,w)\varphi_s(\theta) + (1-\sigma_\theta(q,w))\varphi_l(\theta))) \\
&+ (1-p_h)((\sigma_\theta(q,w)(1-\varphi_s(\theta)) + (1-\sigma_\theta(q,w))(1-\varphi_l(\theta))))V(q,w-1|\theta),
\end{aligned}
$$

with boundary conditions $V(q,0|\theta) = (r-c(q+1))(p_h + (1-p_h)(\sigma_\theta(q,w)\varphi_s(\theta) + (1-\sigma_\theta(q,w))\varphi_l(\theta)))$ for all q. Notice that based on the recursive nature of the value function $V(q,w|\theta)$, the expected social welfare in the system $\Omega(\varphi_s(\theta), \varphi_l(\theta), \theta)$, is given by $V(0, \Lambda - 1|\theta)$. As above, given the fixed threshold structure of the signalling mechanism we have that $\sigma_\theta(q, w) = \sigma_\theta(q)$ for all w, and that $\sigma_\theta(q) = 1$ if $q < \theta$ and $\sigma_\theta(q) = 0$ otherwise.

Estimation of Relevant Parameters. Recall that the QRE of the game is given by $\varphi^*_s(\theta; \beta)$ for $\theta \in \{1, \cdots, \Lambda\}$, $\varphi^*_l(\theta; \beta)$ for $\theta \in \{0, \cdots, \Lambda - 1\}$ and $\vartheta^*(\theta; \beta, \beta_m)$ for $\theta \in \{0, \cdots, \Lambda\}$. Let f denote the data set in this treatment, where we recall that in each round $t \in \{1, \ldots, T\}$ and for each cohort $c \in \{1, \ldots, C\}$, such data set contains the provider j's implemented threshold θ_{jtc}. Moreover, the data set contains for each customer $i \in \{1, \cdots, \Lambda\}$, their strategy $A_{itc} \equiv \{a_{itc}(\varsigma = s|\theta_{jtc}), a_{itc}(\varsigma = l|\theta_{jtc})\}$, where $a_{itc} \in \{0, 1\}$ is such that 1 represents the *Join* action and 0 the *Balk* action. Based on this, the log-likehood function is given by:

$$\begin{aligned}\mathcal{L}(\beta, \beta_m|f) = \sum_i \sum_t \sum_c \big(& a_{itc}(\varsigma = s|\theta_{jtc}) \log(\varphi^*_s(\theta_{jtc}; \beta)) + (1 - a_{itc}(\varsigma = s|\theta_{jtc})) \log(1 - \varphi^*_s(\theta_{jtc}; \beta)) \\ & + a_{itc}(\varsigma = l|\theta_{jtc}) \log(\varphi^*_l(\theta_{jtc}; \beta)) + (1 - a_{itc}(\varsigma = l|\theta_{jtc})) \log(1 - \varphi^*_l(\theta_{jtc}; \beta)) \big) \\ & + \sum_j \sum_t \sum_c \big(\log(\vartheta^*(\theta_{jtc}; \beta, \beta_m)) \big).\end{aligned}$$

We compute the values $\hat{\beta}, \hat{\beta_m}$ that maximize the above function $\mathcal{L}(\beta, \beta_m|f)$. That is, we compute the maximum-likelihood estimates for the parameters β, β_m, based on our experimental data.

Computation of Implied Aggregate Metrics Under Estimated QRE. In Appendix B.2, we describe how to compute the implied values of different aggregate metrics (e.g., expected social welfare Ω, informativeness $\mathcal{I}$) based on joining proportions $\varphi^s_{tc}(\theta_{jtc})$, $\varphi^l_{tc}(\theta_{jtc})$ for given an implemented threshold θ_{jtc}. To compute the implied values of the aggregate metrics under the estimated QRE *for a given implemented threshold* θ, we use the same procedure described in Appendix B.2, but we use instead the the estimated QRE joining proportions $\varphi^*_s(\theta; \hat{\beta})$ and $\varphi^*_l(\theta; \hat{\beta})$. Based on this, if we let $m(\varphi^*_s(\theta; \hat{\beta}), \varphi^*_l(\theta; \hat{\beta}), \theta)$ be an arbitrary aggregate metric m evaluated at a given threshold θ, then the *overall implied value of the metric* (across thresholds θ) is given by $\sum_\theta m(\varphi^*_s(\theta; \hat{\beta}), \varphi^*_l(\theta; \hat{\beta}), \theta) \vartheta^*(\theta; \hat{\beta}, \hat{\beta_m})$.

D.3. No Commitment Treatment

In this case, given a communicated threshold θ' and a signal ς, L type customers join with probability:

$$\varphi_s(\theta'; \beta, \beta_m, \kappa) = \frac{e^{(r - c(\mathbb{E}[Q|\varsigma = s, \theta'] + 1))/\beta}}{1 + e^{(r - c(\mathbb{E}[Q|\varsigma = s, \theta'] + 1))/\beta}} \quad \text{for} \quad \theta' = 0, \cdots, \Lambda, \tag{14}$$

$$\varphi_l(\theta'; \beta, \beta_m, \kappa) = \frac{e^{(r - c(\mathbb{E}[Q|\varsigma = l, \theta'] + 1))/\beta}}{1 + e^{(r - c(\mathbb{E}[Q|\varsigma = l, \theta'] + 1))/\beta}} \quad \text{for} \quad \theta' = 0, \cdots, \Lambda. \tag{15}$$

In these expressions, customers compute the expected numbed of customers in the system $\mathbb{E}[Q|\varsigma,\theta']$ upon their arrival, based on the provider's communicated threshold θ' and their beliefs about how other customers behave. We note that such expectation depends on the providers parameters β_m and κ. The reason is that in this treatment, customers do not observe the implemented threshold θ directly. They can only infer it from their beliefs about the service provider's joint selection of thresholds. Formally, the service provider selects a joint probability distribution $\vartheta(\theta,\theta';\beta,\beta_m,\kappa)$ over the choice of the pair of thresholds θ and θ' (where we abuse notation by dropping the dependence of φ_l and φ_s on β, β_m and κ):

$$\vartheta(\theta,\theta';\beta,\beta_m,\kappa)=\frac{e^{(\Omega(\varphi_s(\theta'),\varphi_l(\theta'),\theta)-\kappa|\theta'-\theta|)/\beta_m}}{\sum_\theta\sum_{\theta'}e^{(\Omega(\varphi_s(\theta'),\varphi_l(\theta'),\theta)-\kappa|\theta'-\theta|)/\beta_m}} \quad \text{for} \quad \theta=0,\cdots,\Lambda \quad \text{and} \quad \theta'=0,\cdots,\Lambda. \tag{16}$$

Based on the above, in equilibrium, customers' and the service provider's expectations and beliefs are consistent with customers' and the service provider's behaviour. That is, the QRE of the game is given by the solution of the system of equations (14) - (16), denoted by $\varphi_s^*(\theta';\beta,\beta_m,\kappa)$ for $\theta'\in\{0,\cdots,\Lambda\}$, $\varphi_l^*(\theta';\beta,\beta_m,\kappa)$ for $\theta'\in\{0,\cdots,\Lambda\}$ and $\vartheta^*(\theta,\theta';\beta,\beta_m,\kappa)$ for $\theta\in\{0,\cdots,\Lambda\}$ and $\theta'\in\{0,\cdots,\Lambda\}$. To compute the QRE, we first need to understand how to compute customers' expectations, beliefs and the service provider's expected social welfare.

How to compute $\mathbb{E}[Q|\varsigma,\theta']$. First, we note that for a fixed implemented threshold θ, we can compute the expected number of customers as follows:

$$\begin{aligned}\mathbb{E}[Q|\varsigma,\theta]&=\sum_{q=0}^{\Lambda-1}q\mathbb{P}(q|\varsigma,\theta)=\sum_{q=0}^{\Lambda-1}q\frac{\mathbb{P}(\varsigma|q,\theta)\mathbb{P}(q|\theta)}{\mathbb{P}(\varsigma|\theta)}\\&=\sum_{q=0}^{\Lambda-1}q\frac{(\sum_{w=0}^{\Lambda-1}\mathbb{P}(\varsigma|q,w,\theta)\mathbb{P}(w|q,\theta))\mathbb{P}(q|\theta)}{\mathbb{P}(\varsigma|\theta)}\\&=\sum_{q=0}^{\Lambda-1}q\frac{\sum_{w=0}^{\Lambda-1}\mathbb{P}(\varsigma|q,w,\theta)\mathbb{P}(q,w|\theta)}{\sum_{w=0}^{\Lambda-1}\sum_{q=0}^{\Lambda-1}\mathbb{P}(\varsigma|q,w,\theta)\mathbb{P}(q,w|\theta)},\end{aligned}$$

where q represents the current number of customers in queue, and w the remaining customers to arrive. Now, since the communicated threshold θ' is not necessarily the same that generates the signals, customers can only infer the implemented threshold θ. It follows that:

$$\mathbb{E}[Q|\varsigma,\theta']=\mathbb{E}\left[\mathbb{E}[Q|\varsigma,\theta]|\theta'\right]=\sum_{\theta=0}^{\Lambda}\mathbb{E}[Q|\varsigma,\theta,\theta']\mathbb{P}(\theta|\theta',\varsigma)=\sum_{\theta=0}^{\Lambda}\mathbb{E}[Q|\varsigma,\theta,\theta']\mathbb{P}(\theta|\theta',\varsigma)=\sum_{\theta=0}^{\Lambda}\mathbb{E}[Q|\varsigma,\theta]\mathbb{P}(\theta|\theta'),$$

where we note that Q is conditionally independent of θ' given θ, and that $\mathbb{P}(\theta|\theta',\varsigma)=\mathbb{P}(\theta|\theta')$ since the provider selects a static signalling policy. To compute the state probabilities $\mathbb{P}(q,w|\theta,\theta')$, we need to consider the dynamics in the queue. Let $\mathbb{P}(\varsigma=s|q,w,\theta)=\sigma_\theta(q,w)$, and $\mathbb{P}(\varsigma=l|q,w,\theta)=1-\sigma_\theta(q,w)$. First, notice that states (q,w) can be reached only from states $(q-1,w+1)$ and $(q,w+1)$. For a given $\varphi_s(\theta')$, $\varphi_l(\theta')$ and θ, if the system is in state $(q-1,w+1)$, then the system transitions into state (q,w) with probability $p_h+(1-p_h)(\sigma_\theta(q-1,w+1)\varphi_s(\theta')+(1-\sigma_\theta(q-1,w+1))\varphi_l(\theta'))$. On the other hand, if the system is in state $(q,w+1)$, then the system transitions into state (q,w) with probability $(1-p_h)(\sigma_\theta(q-1,w+1)(1-\varphi_s(\theta'))+(1-\sigma_\theta(q-1,w+1))(1-\varphi_l(\theta')))$. Based on this, we can capture the queueing dynamics with the following recursive expression:

$$\begin{aligned}\mathbb{P}(q,w|\theta,\theta') =& \mathbb{P}(q-1,w+1|\theta,\theta')(p_h+(1-p_h)(\sigma_\theta(q-1,w+1)\varphi_s(\theta')+(1-\sigma_\theta(q-1,w+1))\varphi_l(\theta')))\\ &+\mathbb{P}(q,w+1|\theta,\theta')(1-p_h)(\sigma_\theta(q-1,w+1)(1-\varphi_s(\theta'))+(1-\sigma_\theta(q-1,w+1))(1-\varphi_l(\theta'))),\end{aligned}$$

with boundary conditions $\mathbb{P}(0,\Lambda-1|\theta,\theta')=1/\Lambda$, $\mathbb{P}(0,w|\theta,\theta')=\mathbb{P}(0,w+1|\theta,\theta')(1-p_h)(\sigma_\theta(0,w+1)(1-\varphi_s(\theta'))+(1-\sigma_\theta(0,w+1))(1-\varphi_l(\theta')))$ for all $w<\Lambda-1$, and $\mathbb{P}(q,w|\theta,\theta')=\mathbb{P}(q-1,w+1|\theta,\theta')(p_h+(1-p_h)(\sigma_\theta(q-1,w+1)\varphi_s(\theta')+(1-\sigma_\theta(q-1,w+1))\varphi_l(\theta')))$ for all $q>0, w<\Lambda-1$ such that $q+w=\Lambda-1$. Note that given the fixed threshold structure of the signalling mechanism we have that $\sigma_\theta(q,w)=\sigma_\theta(q)$ for all w, and that $\sigma_\theta(q)=1$ if $q<\theta$ and $\sigma_\theta(q)=0$ otherwise.

How to compute $\mathbb{P}(\theta|\theta')$ and Ω. To compute customers beliefs about the joint distribution of the service provider $\mathbb{P}(\theta|\theta')=\frac{\vartheta(\theta,\theta')}{\vartheta(\theta')}=\frac{\vartheta(\theta,\theta')}{\sum_\theta\vartheta(\theta,\theta')}$, we consider the value function $V(q,w|\theta,\theta')$. For a given $\varphi_s(\theta')$, $\varphi_l(\theta')$, θ', and θ, in state (q,w), the expected future utility $V(q,w|\theta,\theta')$ is equal to the immediate expected utility, $p_h(r-c(q+1))+(1-p_h)(r-c(q+1))(\sigma_\theta(q,w)\varphi_s(\theta')+(1-\sigma_\theta(q,w))\varphi_l(\theta'))$, plus the expected utility from time $w-1$ onward, $(p_h+(1-p_h)(\sigma_\theta(q,w)\varphi_s(\theta')+(1-\sigma_\theta(q,w))\varphi_l(\theta'))V(q+1,w-1|\theta,\theta')+(1-p_h)((\sigma_\theta(q,w)(1-\varphi_s(\theta'))+(1-\sigma_\theta(q,w))(1-\varphi_l(\theta'))))V(q,w-1|\theta,\theta')$. It follows that:

$$\begin{aligned}V(q,w|\theta,\theta') =& (r-c(q+1)+V(q+1,w-1)|\theta,\theta')(p_h+(1-p_h)(\sigma_\theta(q,w)\varphi_s(\theta')+(1-\sigma_\theta(q,w))\varphi_l(\theta')))\\ &+(1-p_h)((\sigma_\theta(q,w)(1-\varphi_s(\theta'))+(1-\sigma_\theta(q,w))(1-\varphi_l(\theta'))))V(q,w-1|\theta,\theta'),\end{aligned}$$

with boundary conditions $V(q,0|\theta,\theta')=(r-c(q+1))(p_h+(1-p_h)(\sigma_\theta(q,w)\varphi_s(\theta')+(1-\sigma_\theta(q,w))\varphi_l(\theta')))$ for all q. Notice that based on the recursive nature of the value function $V(q,w|\theta,\theta')$, the expected social welfare in the system $\Omega(\varphi_s(\theta'),\varphi_l(\theta'),\theta)$, is given by $V(0,\Lambda-1|\theta,\theta')$. As above, given the fixed threshold structure of the signalling mechanism we have that $\sigma_\theta(q,w)=\sigma_\theta(q)$ for all w, and that $\sigma_\theta(q)=1$ if $q<\theta$ and $\sigma_\theta(q)=0$ otherwise.

Estimation of Relevant Parameters. Recall that the QRE of the game is given by $\varphi_s^*(\theta';\beta,\beta_m,\kappa)$ for $\theta' \in \{0,\cdots,\Lambda\}$, $\varphi_l^*(\theta';\beta,\beta_m,\kappa)$ for $\theta' \in \{0,\cdots,\Lambda\}$ and $\vartheta^*(\theta,\theta';\beta,\beta_m,\kappa)$ for $\theta \in \{0,\cdots,\Lambda\}$ and $\theta' \in \{0,\cdots,\Lambda\}$. Let f denote the data set in this treatment, where we recall that in each round $t \in \{1,\dots,T\}$ and for each cohort $c \in \{1,\dots,C\}$, such data set contains the provider j's implemented threshold θ_{jtc} and communicated threshold θ'_{jtc}. Moreover, the data set contains for each customer $i \in \{1,\cdots,\Lambda\}$, their strategy $A_{itc} \equiv \{a_{itc}(\varsigma = s|\theta'_{jtc}), a_{itc}(\varsigma = l|\theta'_{jtc})\}$, where $a_{itc} \in \{0,1\}$ is such that 1 represents the *Join* action and 0 the *Balk* action. Based on this, the log-likehood function is given by:

$$\begin{aligned}\mathcal{L}(\beta,\beta_m,\kappa|f) = \sum_i\sum_t\sum_c \big(&a_{itc}(\varsigma = s|\theta'_{jtc})\log(\varphi_s^*(\theta'_{jtc};\beta,\beta_m,\kappa)) + (1-a_{itc}(\varsigma = s|\theta'_{jtc}))\log(1-\varphi_s^*(\theta'_{jtc};\beta,\beta_m,\kappa)) \\ &+ a_{itc}(\varsigma = l|\theta'_{jtc})\log(\varphi_l^*(\theta'_{jtc};\beta,\beta_m,\kappa)) + (1-a_{itc}(\varsigma = l|\theta'_{jtc}))\log(1-\varphi_l^*(\theta'_{jtc};\beta,\beta_m,\kappa))\big) \\ &+ \sum_j\sum_t\sum_c \big(\log(\vartheta^*(\theta_{jtc},\theta'_{jtc};\beta,\beta_m,\kappa))\big).\end{aligned}$$

We compute the values $\hat{\beta},\hat{\beta_m},\hat{\kappa}$ that maximize the above function $\mathcal{L}(\beta,\beta_m,\kappa|f)$. That is, we compute the maximum-likelihood estimates for the parameters β,β_m,κ, based on our experimental data.

Computation of Implied Aggregate Metrics Under Estimated QRE. In Appendix B.2, we describe how to compute the implied values of different aggregate metrics (e.g., expected social welfare Ω, informativeness $\mathcal{I}$) based on joining proportions $\varphi_{tc}^s(\theta'_{jtc})$, $\varphi_{tc}^l(\theta'_{jtc})$ for given pair of communicated θ'_{jtc} an implemented thresholds θ_{jtc}. To compute the implied values of the aggregate metrics under the estimated QRE *for a given pair of thresholds* θ' and θ, we use the same procedure described in Appendix B.2, but we use instead the the estimated QRE joining proportions $\varphi_s^*(\theta';\hat{\beta},\hat{\beta_m},\hat{\kappa})$ and $\varphi_l^*(\theta';\hat{\beta},\hat{\beta_m},\hat{\kappa})$. Based on this, if we let $m(\varphi_s^*(\theta';\hat{\beta},\hat{\beta_m},\hat{\kappa}),\varphi_l^*(\theta';\hat{\beta},\hat{\beta_m},\hat{\kappa}),\theta',\theta)$ be an arbitrary aggregate metric m evaluated at a pair of thresholds θ' and θ, then the *overall implied value of the metric* (across thresholds θ) is given by $\sum_{\theta'}\sum_{\theta} m(\varphi_s^*(\theta';\hat{\beta},\hat{\beta_m},\hat{\kappa}),\varphi_l^*(\theta';\hat{\beta},\hat{\beta_m},\hat{\kappa}),\theta',\theta)\vartheta^*(\theta',\theta;\hat{\beta},\hat{\beta_m},\hat{\kappa})$.