

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

Credibility and Effectiveness of Information Design in Service Operations

Arturo Estrada Rodriguez

Católica Lisbon School of Business and Economics, Universidade Católica Portuguesa, 1649-023 Lisbon
aerodriguez@ucp.pt

Rouba Ibrahim

School of Management, University College London, 1 Canada Square, London E14 5AB
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 information sharing as a managerial lever to mitigate customer overuse of resources in service systems. We develop a queuing communication game and report results from laboratory experiments that test core predictions of information design theory in service operations. We find that information design increases social welfare not only through mechanisms of *Bayesian persuasion*—when providers can credibly commit to their information policy—but also in *Cheap talk* settings without formal commitment, with the greatest welfare gains occurring when providers strategically overstate queue lengths. Our findings shed light on the extent to which formal commitment mechanisms and behavioral mechanisms can function as substitutes—challenging the necessity of costly commitment devices. We discuss how this behavioral perspective broadens the scope of effective information policies.

Key words: information sharing; commitment; persuasion; cheap talk; queueing; noisy decision-making

1 Introduction

A well-established problem in the service operations literature is the tendency of self-interested customers to join congestion-prone systems at higher-than-socially optimal rates (Naor 1969, Chen and Hasenbein 2020, Hassin and Haviv 2003, Hassin 2016). *Overjoining* is an ubiquitous issue in service settings, including healthcare, transportation, customer support, or even theme parks. In many cases, excessive demand strains limited resources, leading to longer wait times, reduced service quality, and increased operational costs. Importantly, problems compound when overjoining behavior imposes negative externalities on customers with more pressing needs. For example, in the Portuguese healthcare system, patients with low medical urgency overcrowd emergency rooms, limiting access to resources for critical cases (Lusa 2023). Similarly, in the UK, nonemergency ambulance calls extend wait times for genuine emergencies (Bramwell 2023).

Addressing overjoining—especially by low-need customers—is challenging. Capacity expansion (e.g., more staff or faster processing) is often prohibitively costly. Strict admission controls can be unfair or infeasible (Haviv and Oz 2018), and regulations in many sectors limit direct control over access. For instance, in UK healthcare, patients have the legal right to choose their outpatient provider (NHS 2023). Monetary incentives such as entry fees (Naor 1969, Edelson and Hilderbrand 1975), entry contracts (Haviv 2014), or pay-for-priority mechanisms (Mendelson and Whang 1990, Hassin 1995) are also inapplicable in publicly funded systems like healthcare or social housing. For example, European regulatory frameworks impose public service obligations that explicitly prioritize need-based access over ability to pay (EU-Commission 2020). However, while prioritization strategies can help, accurately assessing unobservable customer service needs is challenging without added resource strain (Argon and Ziya 2009, Rodriguez et al. 2025). Under these conditions, *information design*—selectively sharing system information to influence customer behavior—is an attractive managerial lever to control demand for scarce resources (Bergemann and Morris 2019).

Aimed at deterring demand from low-need customers for the benefit of those with higher service needs, information design is a cost-effective non-intrusive mechanism that can leverage advances in communication technologies. For example, UK healthcare systems are experimenting with mobile apps showing real-time delays at A&E departments, minor injury units, and walk-in centers, aiming to divert low-urgency cases and reduce delays for urgent ones (GovUk 2024). However, information design presents service providers with a nontrivial challenge. Because the provision of full information is 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 information to have the desired effect on customer behavior, customers need to find it credible. What then are the conditions under which service providers can effectively use information to address the overjoining problem in service systems? We can draw initial insights from existing literature.

A large body of theoretical work explores how rational customers respond to different levels of delay information, from full to partial to none (see Ibrahim 2018 and Economou 2021 for reviews). A key assumption in this literature is that providers can *commit* to an information policy, that is, that the communicated information truthfully reflects system states, which ensures credibility that allows customers to trust the information and act accordingly. In contrast, without commitment, rational customers anticipate misaligned incentives and may completely disregard information, rendering communication ineffective (Allon et al. 2011). Thus, effective information design, in practice, relies on costly commitment devices. These can include waiting-time guarantees and compensation (Kumar et al. 1997), publicly scrutinized internet platforms (Singh et al. 2023), wait-time-based pricing (Lin et al. 2023), blockchain-based protocols (Drakopoulos et al. 2022, Narayanan 2023), or audit trails. A

UK Supreme Court case, where hospital staff were held liable for misleading wait-time information, illustrates the importance of such trails (Hyde 2018).

Potentially softening the need for costly commitment mechanisms are several results from behavioral economics which suggests that credible communication can occur even when standard economic theory predicts uninformative “cheap talk” (Blume et al. 2020). However, it is less obvious whether and how these results from relatively generic settings translate to operationally richer service settings which feature multiple customers (i.e., receivers of information) with endogenous payoff dependencies between customer choices and dynamically evolving system states. Moreover, the existing behavioral literature leaves open the question of whether behavioral mechanisms can genuinely substitute for formal commitment. Indeed, there appear to be limits to credibility in service settings without an obvious commitment mechanism. For example, Disney World visitors have raised concerns about inflated wait-time announcements at popular rides, possibly to divert customers to less popular attractions and help the park absorb demand (Bricker 2024). Similarly, customer-service departments across various industries—including banking, broadband, mobile, credit cards, energy, and water—have been criticized for routinely using the message “we are experiencing high call volumes” regardless of the actual system state, as a tactic to encourage more customers to access digital and self-service channels (Lazarus 2020, Fisher 2025).

Our study provides first empirical answers to whether credible and effective communication can arise *in equilibrium* between a service provider and a customer population, and what role commitment plays in this interaction. Robustly answering this is challenging. It demands significant environmental control (e.g., of customer needs and the provider’s ability to commit), observation of provider choices, and a theoretical benchmark to gauge communication effectiveness. Since field settings rarely meet these conditions, we instead develop a queuing communication game that provides sharp theoretical predictions, which we subsequently test under controlled laboratory conditions where we vary the (human) provider’s ability to communicate information on queue status to their (human) customers.

Our model is ideally suited for behavioral queuing experiments, and it accurately replicates key operational features of real service systems in a lab setting. These features—like customer heterogeneity, negative externalities, and dynamic queue formation—distinguish service systems from other areas where strategic communication has been studied (e.g., Fréchet et al. 2022). In our model, customers, who arrive in random order and differ in service need, decide whether to join a queue or opt for an outside alternative. We characterize the conditions under which rational, self-interested, low-need customers join the system, disregarding the delays they impose on subsequent arrivals. A conflict of interest arises when a service provider aims to maximize social welfare by diverting low-need demand to shorten waiting times for high-need customers. In baseline *No-Information* and *Full-Information* settings, low-need customers overjoin. We show that sharing partial queue information (a “short”

or “long” wait message) can deter low-need customers from joining, but this relies on credibility. In a *Commitment* setting, providers must share truthful information, ensuring credibility and leading to welfare gains through *Bayesian persuasion*. Here, providers send intentionally ambiguous but truthful signals that influence low-need customer behavior. However, in a *No-Commitment* setting, providers can misreport—for example, by inflating delay announcements. This leads to a breakdown in communication, and welfare reverts to the *No-Information* baseline as rational low-need customers disregard untrustworthy signals.

We conduct multiple experiments that vary information “regime” to test key theoretical predictions. To study behavioral mechanisms and boundary conditions, our experiments systematically vary population size and the salience of conflicting incentives between customers and providers. In our experiments, customers approach random behavior in *No-Information* treatments, and follow signals when service providers share information, even in the absence of a commitment mechanism. Providers send informative signals in all treatments, regardless of their ability to formally commit to their information policy. Importantly, they rarely use ambiguous signals in the *Commitment* treatment, but often strategically inflate wait signals in the *No-Commitment* treatment—consistent with anecdotal evidence from different service industries (Bricker 2024, Lazarus 2020, Fisher 2025). We observe that information design improves welfare due to the interplay between customer and provider behavior. However, this is not entirely consistent with standard theory; the greatest welfare improvements occur in the absence of a commitment mechanism, when providers lie by inflating wait-time announcements to deter low-need customers. To quantitatively reproduce the main departures between our data and theoretical predictions, we estimate a structural model of quantal response equilibrium (QRE). This model parsimoniously captures key behavioral patterns in our data, in particular customers’ limited ability to perfectly infer queue states from available information (Huang et al. 2013).

Our findings offer several important managerial implications. Generally, we show that carefully designed information policies allow firms to mitigate the overjoining problem that plagues many service systems in practice. Importantly, our results shed light on the extent to which formal commitment mechanisms and behavioral mechanisms can be viewed as substitutes. When formal commitment devices are prohibitively expensive (or simply not available), service providers can improve social welfare with carefully inflated delay announcements. When a formal commitment device is available, although cognitively challenging to implement for human players, providers can mitigate overjoining behavior by using informative obfuscation—deliberately vague or pooled signals.¹ Relatedly, our study highlights the benefits of a limited message space (e.g., “short” vs. “long” delays), which not only simplifies implementation but also leads to better outcomes than fully revealing the queue state.

¹ The direct evidence for this claim stems from an additional experimental treatment that automates the service provider, and implements the optimal (“persuasive”) policy, and thus controls for human provider bias.

The remainder of the paper is structured as follows. We review the relevant literature in §2, develop the theoretical model in §3, and describe our experiments and results in §4. We discuss in §5 a behavioral account that accommodates the main patterns in the data, and discuss various managerial insights that arise from our results. In §6 we conclude the paper.

2 Related Literature

Our work relates to several streams of theoretical and empirical research in economics and operations.

The seminal works of Naor (1969) and Edelson and Hilderbrand (1975) pioneered the study of how queue observability affects customer behavior. Since then, researchers have investigated whether fully revealing or fully concealing the queue improves system-level outcomes (Hassin 1986, Chen and Frank 2004, Shone et al. 2013, Chen and Hasenbein 2020). The general insight is that neither extreme is uniformly preferable across all operational and economic parameters. This has established information sharing as an extensively studied managerial lever, with the literature exploring various partial information structures; see Ibrahim (2018) and Economou (2021) for comprehensive surveys. More recently, researchers have applied *Bayesian Persuasion* frameworks (Kamenica and Gentzkow 2011, Bergemann and Morris 2019) to show that strategic obfuscation can induce desirable customer behavior (Lingenbrink and Iyer 2019, Che and Tercieux 2021).

The overarching goal of our study is to provide a first empirical test of whether effective communication can arise in equilibrium between a service provider and a population of customers that arrive at a congestion-prone system. 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. A central assumption in their study, common in theoretical work on information design, is that providers can commit to the information they announce—ensuring it is truthful and credible. With commitment, customers trust and act on the message. Without it, rational customers anticipate misreporting, disregard the information, and information sharing can become ineffective cheap talk (Allon et al. 2011). Our setting shares with the theoretical analysis of Anunrojwong et al. (2022) key features of the queuing environment (with important differences that we discuss in §3), but our objective is different. Specifically, we ask: Do the theoretical benefits of persuasive communication survive a test with human service providers and customers, and what is the role of commitment?

We can look for initial answers to these empirical questions in two related streams of behavioral work on *Bayesian Persuasion* and on *Cheap Talk*. Following the foundational theoretical work of

Kamenica and Gentzkow (2011), a fairly recent experimental literature has begun to test behavioral mechanisms and boundary conditions of the Bayesian persuasion framework. The results so far are mixed. Nguyen (2017) observe a high fraction of senders and receivers converge to the optimal persuasive strategy, and Coricelli et al. (2023) find that senders are more successful in influencing receivers via persuasion than via monetary incentives. In contrast, Au et al. (2023) document a high rate of persuasion failure, and their results suggests that senders have difficulties implementing optimal strategies when facing high strategic uncertainty in receiver behavior - a result with likely relevance for a queuing setting with substantial customer decision-making noise. While models and experiments of Bayesian persuasion endow the sender with the ability to communicate according to publicly known predetermined protocols, a well established behavioral literature on cheap talk settings suggests that information sharing can be effective even in the absence of a formal commitment mechanism (e.g., Cai and Wang 2006, Kawagoe and Takizawa 2009, Wang et al. 2010, Blume et al. 2020). The cheap talk literature has systematically documented over-transmission of information, where a substantial proportion of players tend to be truthful as a sender and credulous as a receiver. Although the behavioral literature suggests that effective communication can arise when credibility of information is ensured by a commitment mechanism, or by behavioral mechanisms in the absence of commitment, it largely leaves open the questions about the conditions under which one would be preferred over the other. To our knowledge, Fr chet te et al. (2022) is the only experimental study that, like us, explicitly varies commitment in a unified framework-and they find that commitment improves outcomes as theory predicts.

The experimental economics literature provides several insights on information design, and the role of commitment in it, but it is not a priori obvious what insights translate to operationally richer service settings that we study. Fr chet te et al. (2022) is a case in point—in contrast to their main finding, we find that commitment yields little benefit in our operational queueing setting. And the effect can even be directionally opposite: the most pronounced benefits arise when providers strategically misreport queue lengths in a setting without commitment. We attribute this divergence to the greater complexity of our queueing setting, which reflects realistic operational dynamics and places greater cognitive demands on participants. Unlike static single-sender-single-receiver games typically studied in the economics literature, providers interact with a sequence of customers about dynamically evolving queue states. This creates interdependent decisions and strategic externalities—early actions shape later beliefs. Moreover, rather than exogenous conflicts of interest, the incentive conflict between provider and customers in our setting results from fundamental queueing dynamics. This underscores the importance of testing foundational economic theories in operational contexts.

Our focus on the credibility of operational information also relates to a broad and growing Behavioral Operations literature. One stream of work documents the (mostly positive) effects of information sharing on various customer-level outcomes in service settings—including perceived value and

willingness-to-pay (Buell and Norton 2011, Buell et al. 2017), trust and engagement in co-production (Buell et al. 2021), ex-post perceptions of wait times and service evaluations (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), time spent in service (Webb et al. 2020), 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 misalignment (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.

The main results from our study are explicable by key behavioral concepts from economics that have recently gained popularity in the Behavioral Operations literature. Specifically, like us, several previous studies document deviations from uninformative equilibrium predictions in operations settings due to *other-regarding preferences* such as trust, trustworthiness, or lying aversion. Unlike these studies, we explicitly design our experiments to address the role of commitment, and explore the managerial implications for information design, rather than forecast policies (Özer et al. 2011, Özer et al. 2018), incentive design (Scheele et al. 2018), or scheduling practices (Rodriguez et al. 2025). Similarly, like us, several previous studies address noisy decision making, often discussed in the operations literature under the term *bounded rationality* (Su 2008, Chen et al. 2012, Kremer and Debo 2016, Goldschmidt et al. 2021). Close to our setting is the theoretical work of Huang et al. (2013), who study the impact of bounded rationality on decision making in observable and unobservable queue settings through the lens of the quantal response equilibrium (QRE) framework. Intriguingly, the authors highlight conditions under which bounded rationality can improve social welfare. Our setting allows us to empirically test some of their key implications related to information availability.²

Finally, while our general focus is on service systems, this work provides insights 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), retail

² Note that models of bounded rationality (such as QRE or level-K) are often used to organize behavioral patterns observed in Bayesian persuasion and Cheap talk settings (e.g., Wu and Ye 2023, Fréchette et al. 2022, Cai and Wang 2006, Kawagoe and Takizawa 2009, Wang et al. 2010).

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 Modelling Framework

We study how service providers can influence customer joining decisions by sharing information, and how the amount of information provided affects social welfare. We develop a model that shares with Anunrojwong et al. (2022) several characteristics of service systems, as well as some of its key predictions, but our model is designed specifically for experimental testing. Like Kremer and Debo (2016) and Frazelle and Katok (2024), we depart from the standard queueing model framework, which assumes stochastic arrival and service processes, and focuses on steady-state analysis. Specifically, we consider a system with deterministic service times, where a finite number of customers make sequential joining decisions in random order, and service begins only after all customers have made their decisions. By simplifying our model, we can faithfully reproduce it in a laboratory setting. This reduces noise in the experimental data and simplifies the setup for participants.

Crucially, our model retains the features most salient for, and specific to, decision-making in queueing environment. First, the queue length, which captures the state of the system, is *dynamically driven* by customer sequential joining decisions. Second, each customer makes decisions based on *expected* queue lengths, and decisions are made before any service begins. Third, customers exert *negative externalities* since joining decisions increase wait-related costs for later arrivals. Finally, our model captures the service provider’s *informational advantage* over customers. Thus, the service provider can influence customer behavior by disclosing information about the state of the system.

3.1 Queueing Model

Λ customers arrive at the system, each with a randomly assigned index $k \in \{1, \dots, \Lambda\}$. Customers make joining decisions sequentially, in increasing order of their assigned index. There is a single server which starts processing customers in a first-come-first-serve order once all customers have made their joining decisions. Customers do not know their indices, nor can they observe the indices of other customers, but they know all other system parameters. For experimental amenability, we assume that service times are deterministic and normalized to 1. Similar to the case with stochastic service times, in our setting customers make joining decisions based on the expected queue length, and service begins after decisions are made.

We consider two types of customers $i \in \{H, L\} = \{\text{high-need}, \text{low-need}\}$ who receive *reward* r_i from service, where $r_H \geq r_L$. Customers are high-need with probability p_H or low-need with probability $1 - p_H$. The service provider cannot observe the types of individual customers, but knows the probability distribution of customer types. We let U_k denote the utility accrued by customer k , and assume

that customers make decisions to maximize their expected utility $\mathbb{E}[U_k]$. The realized utility that customers accrue by joining the queue at the $(q + 1)^{st}$ position, i.e., with q customers ahead, is $u_k = r_i - c(q + 1)$. That is, customers earn a reward r_i upon completion of their service, and incur a delay cost $c(q + 1)$, where c represents customer delay sensitivity. We let v_i denote the expected utility that customers of type i experience from the outside option, that is, if they decide to balk. We assume that the reward r_i is sufficiently high for both types, $r_i - c > v_i$, making it possible to choose joining the system over the outside option. We assume that $r_L - c\Lambda < v_L$, ensuring that low-need customers prefer the outside option if the queue is sufficiently long. Finally, we assume that $r_H - c\Lambda > v_H$, ensuring that high-need customers always prefer to seek service from the system over the outside option; see Anunrojwong et al. (2022) for a similar assumption. As in that paper, this leads to a model that satisfies *utility dominance*, meaning that the *incremental utility*—the net benefit over the outside option—for high-need customers, given by $r_H - c(q + 1) - v_H$, is greater than that for low-need customers for all q . That is, for any queue length, high-need customers derive greater benefit from joining the system than low-need customers do. In §3.5.2, we discuss the importance of customer heterogeneity for the study of information design in queueing models.

We study how the provider can share queue-length information to influence customer behavior and improve expected social welfare $\Omega \triangleq \mathbb{E} \left[\sum_{k=1}^{\Lambda} U_k \right]$. Specifically, in a setting that is characterized by a fundamental information asymmetry (about queue length) between service provider and customers, we study the impact of three information-sharing structures:

- (a) *No-Information*: Arriving customers do not have information on queue length.
- (b) *Full-Information*: Arriving customers have complete information about queue length.
- (c) *Partial-Information*: Arriving customers receive a binary queue length signal (e.g., long/short).

Given the aforementioned parameter space, we note that high-need customers always join with probability 1 in equilibrium, regardless of the information structure. Based on this, we focus in what follows on how information (or its absence) shapes the behavior of low-need customers.

3.2 No Information

We first study the case where customers do not have information about queue length. In this case, a low-need customer joins with probability $\alpha \in [0, 1]$. To make a joining decision, the customer computes the expected utility of joining $r_L - c(\mathbb{E}[Q] + 1)$. Here, $\mathbb{E}[Q]$ is calculated based on system parameters and the customer's belief about how other customers behave. In Proposition 1, we quantify the equilibrium joining probability α^* for self-interested low-need customers, and the socially optimal joining probability α_{soc}^* . We also compare α^* and α_{soc}^* .

PROPOSITION 1. (*No-Information*). *We let $r'_i \triangleq r_i - v_i$. For low-need customers,*

(a) A unique customer joining probability $\alpha^* \in [0, 1]$ exists, and is given by

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

(b) A unique socially optimal probability $\alpha_{soc}^* \in [0, 1)$ exists, and is given by

$$\alpha_{soc}^* = \begin{cases} 0 & \text{if } \frac{r'_L}{c} \leq 1 + p_H(\Lambda - 1), \\ \frac{r'_L - c(1+p_H(\Lambda-1))}{c(\Lambda-1)(1-p_H)} < 1 & \text{if } \frac{r'_L}{c} > 1 + p_H(\Lambda - 1). \end{cases}$$

(c) Low-need customers overjoin the system, $\alpha_{soc}^* \leq \alpha^*$, and strict inequality holds whenever $\alpha^* > 0$.

3.3 Full Information

With *Full-Information*, customers know the length of the queue, q , upon arrival. In this case, a low-need customer's joining strategy is characterized by a probability vector $\alpha_q \triangleq (\alpha_0, \alpha_1, \dots, \alpha_{\Lambda-1})$, where $\alpha_q \in [0, 1]$ is the joining probability conditional on finding q customers in the system. To make a joining decision, customers calculate the expected utility, $r_L - c(q+1)$, from joining in position $q+1$. Conditioning on the queue length information eliminates the need to infer others' behavior, unlike the *No-Information* case. In Proposition 2, we characterize the equilibrium joining probability vector α_q^* for self-interested low-need customers, and the socially optimal probability vector, $\alpha_{q,soc}^*$.

PROPOSITION 2. (Full-Information). *For low-need customers in the Full-Information setting,*

(a) A unique customer equilibrium threshold joining strategy exists, with threshold $q^* = \lfloor \frac{r'_L}{c} \rfloor$. That is, $\alpha_q^* = 1$ for $q < q^*$ and $\alpha_q^* = 0$ for $q \geq q^*$.

(b) A unique socially optimal threshold joining strategy exists, with threshold q_{soc}^* given by

$$q_{soc}^* = \begin{cases} \lfloor \frac{r'_L - c\Lambda p_H}{c(1-p_H)} \rfloor & \text{if } p_H \leq \frac{r'_L}{c\Lambda}; \\ 0 & \text{if } p_H > \frac{r'_L}{c\Lambda}. \end{cases}$$

(c) Low-need customers overjoin the system, $q_{soc}^* \leq q^*$, and strict inequality holds for $p_H > r'_L/c\Lambda$.

Note that the *Full-Information* case allows for two interpretations in our setting, each with their own challenges. First, Proposition 2(a) corresponds to the well-studied case of an *observable queue*, which removes information asymmetry about queue length and thus relieves the provider from carefully designing communication about queue length and credibility concerns. However, rendering a queue (literally) observable is not possible in many practical settings, and it might in of itself create conditions for overjoining behavior (Proposition 2(c))³. Second, full information could describe

³ In contrast to the *M/M/1* setting—where welfare losses arise even with homogeneous customers—Proposition 2 shows that in our model, selfish and socially optimal behavior align when $p_H = 0$. For a detailed comparison between the two modeling frameworks and the reason for such difference, we refer the interested reader to Appendix B.

the *unobservable queue* case in which a better informed service provider communicates to arriving customer the exact queue length—a scenario where conflicts of interest raise important credibility concerns. This adds complexity to both the theoretical analysis and the experimental implementation. With these practical challenges in mind, our analysis includes *Full-Information* (observable queue) as a useful benchmark for the simpler *Partial-Information* policy that is often observed in practice, and that is the focus of our study.

3.4 Partial Information

Propositions 1 and 2 replicate in our transient queue model some known results that demonstrate, in the classic $M/M/1$ queue model, socially suboptimal (over)joining for the *No-Information* (Edelson and Hilderbrand 1975) and the *Full-Information* settings (Naor 1969). Our model also predicts overjoining when rational and self-interested customers ignore the impact of their decisions on other customers, with a nuanced difference in the nature of the negative externality. In our model, welfare suffers specifically because low-need customers overcrowd the system at the detriment of high-need customers (we discuss the role of customer heterogeneity in §3.5.2).

Because overjoining under both *No-Information* and *Full-Information* reduces social welfare, the service provider can benefit from strategically sharing partial information to discourage low-need customers from joining. To derive sharp theoretical predictions and experimentally test the effect of information design, we focus on the case of maximal overjoining: $\alpha^* = 1$, $\alpha_{soc}^* = 0$, and $q_{soc}^* = 0$. This implies that the provider prefers to serve only high-need customers. Based on Propositions 1 and 2, this occurs when $\frac{r'_L}{c} \geq \frac{\Lambda-1}{2} + 1$ and $\frac{r'_L}{c} \leq \Lambda p_H$, requiring $p_H \geq \frac{\Lambda+1}{2\Lambda}$.

The service provider sends a binary signal $\varsigma \in \{s, l\} = \{\text{short}, \text{long}\}$, e.g., either short or long wait, to each arriving customer. The service provider *commits* to a signaling mechanism with threshold θ : A customer receives a short signal if they arrive when there are fewer than θ customers in the system, and a long signal otherwise. Optimal binary signaling mechanisms in $M/M/1$ models, which exhibit a threshold structure, have been identified in previous work (Lingenbrink and Iyer 2019, Anunrojwong et al. 2022). While these optimal mechanisms often involve randomized signals at the threshold, we avoid randomization to enhance clarity and understanding in our experiments, while still capturing the key managerial insights of such models.

Although customers cannot directly observe the system state q , they can infer the expected queue length based on the signals provided and their beliefs about the strategies of other customers. A low-need customer's joining strategy is characterized by the pair of conditional joining probabilities $\alpha_s \triangleq \mathbb{P}(\text{Join}|\varsigma = \text{short})$ and $\alpha_l \triangleq \mathbb{P}(\text{Join}|\varsigma = \text{long})$, given signals ς . To build intuition, in Lemma 1 we analyze the customer equilibrium (α_s^*, α_l^*) that arises for a given signaling threshold θ .

LEMMA 1. (**Customer equilibrium for fixed θ**) For a given signaling threshold $\theta \in \{0, \dots, \Lambda\}$, a unique customer equilibrium (α_s^*, α_l^*) exists, given by $\alpha_s^* = 1$ and

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

where $\bar{\theta} = \lfloor \frac{2r'_L - c(\Lambda+1)}{c} \rfloor$, $\bar{\bar{\theta}} = \lceil \frac{2r'_L - 2c - cp_H(\Lambda-1)}{c(2-p_H)} \rceil$, and $0 \leq \bar{\theta} < \bar{\bar{\theta}} < q^*$.

Lemma 1 shows that low-need customers always join after short signals, but become less likely to join after long signals as θ increases—eventually stopping altogether for $\theta \geq \bar{\bar{\theta}}$ ($\alpha_l^* = 0$). This is because long signals increasingly indicate longer waits. Lemma 1 also shows that the service provider can implement thresholds as low as $\theta = \bar{\theta}$ while maintaining the desired outcome: low-need customers who receive long signals choose not to join ($\alpha_l^* = 0$). The lower the threshold the more low-need customers the provider is able to deter from joining and thus to maximize social welfare. This insight is formalized in Proposition 3⁴ and illustrated in Figure 1a using the parameters of some of our experiments.


PROPOSITION 3. (**Optimal persuasive signaling mechanism**) Let $g(\bar{\theta}) \triangleq 2(r'_L - c\bar{\theta}) + (\Lambda - \bar{\theta})p_H(2r'_L - c(2(1 + \bar{\theta}) + p_H(\Lambda - \bar{\theta} - 1)))$. There is a unique equilibrium between the service provider and customers. If $g(\bar{\theta}) \geq 0$, then $\theta^* = \bar{\theta}$, $\alpha_s^* = 1$, and $\alpha_l^* = 0$. Otherwise, $\theta^* = \bar{\theta} - 1$, $\alpha_s^* = 1$, and $\alpha_l^* \in (0, 1)$.

3.5 Informativeness, Persuasion, and Social Welfare

With a binary signaling mechanism, the provider can induce the same customer behavior—and thus achieve the same social welfare—as in *No-Information* by selecting $\theta \in \{0, \Lambda\}$. For $\theta = 0$ ($\theta = \Lambda$), customers receive only long (short) signals. Because these signals are uninformative about the system state, low-need customers always join with probability $\alpha_l^* = 1$ ($\alpha_s^* = 1$). Similarly, by setting $\theta = q^*$, the provider can replicate *Full-Information* outcomes. Low-need customers join with certainty on short signals ($\alpha_s^* = 1$) and balk on long ones ($\alpha_l^* = 0$), effectively implementing the q^* -threshold strategy and associated social welfare. Importantly, Proposition 3 shows that neither extreme—*No Information* ($\theta \in \{0, \Lambda\}$) nor *Full Information* ($\theta = q^*$)—is optimal. Instead, optimal signaling involves *partial* information, with a threshold $0 < \theta^* < q^*$. This choice obfuscates queue states when sending a long signal by pooling favorable states ($q < q^*$, where joining yields higher utility) with unfavorable ones

⁴ The literature typically characterizes an obedient equilibrium where signals are followed (Anunrojwong et al. 2022, Lingenbrink and Iyer 2019). This characterization is without loss of generality when the information designer has access to all possible signaling mechanisms (Bergemann and Morris 2019). However, to facilitate our experiments, we restricted our analysis to a fixed threshold signaling mechanism. Under this constraint, assuming an obedient equilibrium ($\alpha_s^* = 1$ and $\alpha_l^* = 0$) is not necessarily without loss of generality, as demonstrated in Proposition 3.

Figure 1 Social Welfare and Customer Behavior ($\Lambda = 4$, $r_H = 300$, $r_L = 200$, $c = 50$, $v_H = v_L = 75$, $p_H = 0.7$).



($q \geq q^*$, where joining yields lower utility). In this way, long signals deter customer joining via *persuasion* because customers choice of balking means that sometimes they will balk in states where joining would have been better ($\theta^* \leq q < q^*$).

3.5.1 The Role of Commitment. The above analysis relies on the assumption that the service provider *commits* to a signaling mechanism with threshold θ . With commitment, the service provider adheres to the announced threshold. This guarantees credibility in communication. In contrast, without commitment, the provider may deviate from the communicated threshold. In this case, signals become mere “cheap talk” with no informational value, as they may not reflect the true queue states. Credibility then depends on whether the provider has an incentive to deviate from the communicated threshold in favor of an alternative signaling policy. Figure 1b shows that when low-need customers respond to long wait signals by balking, the provider has a clear incentive to inflate delay announcements—sending only long signals (i.e., $\theta = 0$)—in order to discourage low-need customers from joining the queue. Consequently, rational customers anticipate this misalignment of incentives and disregard the messages, effectively reducing communication to the *No-Information* case. Proposition 4 formalizes this result.

PROPOSITION 4. (No Commitment) *Without commitment, only uninformative equilibria exist where the provider implements any $\theta^* \in \{0, \dots, \bar{\theta}, \Lambda\}$ randomly and customers join with probabilities $\alpha_s^* = 1$ and $\alpha_l^* = 1$.*

This means that commitment renders the customer-provider conflict of interest irrelevant, which allows for the possibility of sustaining a credible, informative, and welfare-improving equilibrium.

3.5.2 The Role of Customer Heterogeneity. In our setting, the presence of high-need customers allows for the *persuasion* of low-need customers: upon receiving a long signal, low-need customers may decide to balk, knowing that the queue might be in a favorable state to join, specifically $\theta^* \leq q < q^*$. When customers are homogeneous ($p_H = 0$), such a persuasive threshold ($\theta^* < q^*$) cannot be sustained in equilibrium. To see why, suppose instead that $\theta^* < q^*$ holds in equilibrium with homogeneous customers ($p_H = 0$) who balk upon receiving a long signal. Under this assumption, a customer receiving a long signal would know with certainty that the queue is shorter than q^* , and would therefore strictly prefer to join—a contradiction. Thus, persuasion is not possible with a homogeneous population. In contrast, with high-need customers who join, the queue may exceed q^* . This allows a low-need customer who receives a long signal to be deterred from joining. See §4.1 in Anunrojwong et al. (2022) for a related discussion in the $M/M/1$ setting.

3.6 Numerical Example: Information Design and Welfare

We illustrate how information design improves welfare in our queuing environment for a small population of $\Lambda = 4$ customers, using the parameters from experimental Studies 1 and 2. For each of 16 unique arrival sequences (i_k denotes a customer of type $i \in \{L, H\}$ in arrival position k), Table 1 displays queues and customer utilities, as they evolve from the interplay of arrival sequence and equilibrium behavior under each of the information environments. Importantly, the table highlights when an individual customer benefits (i_k) or not ($\bar{i}_k$), from their own or other customers' choices.

In the *No-Information* baseline (Proposition 1), rational customers always join under the chosen parameters. Information design can increase welfare by inducing low-need customers to balk, as this may benefit later arriving high-need customers and sometimes the low-need customers themselves.

In the *Full-Information* setting (Proposition 2), observing the queue length allows low-need customers to balk at high congestion levels when they would otherwise have erroneously joined in *No-Information* (L_3 and L_4), and welfare improves further when a later arriving high-need customer faces a less congested system (H_4). Together, the two effects increase welfare Ω from 580 to 620.5.

The *Partial-Information* setting provides an opportunity to improve welfare further, by persuading low-need customer to balk even when it is not in their own best interest. Although the low-need customer's utility reduces by 75 (L_1) or 25 (L_2) from balking at low congestion levels where they would join in *Full-Information*, welfare improves when there is a sufficiently high probability of later arriving high-need customers (H_k) each of which gains $c = 50$ from every previous low-need customer that balked. The table illustrates the welfare benefits and limits to this approach.

With access to a *Commitment* device (Proposition 3), the provider sends a “short wait” message only when the queue is empty (i.e., a threshold $\theta^* = 1$), and *persuades* some low-need customers to balk when they receive the “long wait” signal—as it effectively pools unfavorable queue positions at

which low-need customers prefer to balk (L_3, L_4) with favorable positions at which they would prefer to join (L_2). Such a properly designed signaling policy improves welfare to $\Omega = 634$.

Notably, when customers *obediently* follow signals, a provider *without commitment* has an incentive to send only “long wait” messages (i.e., a threshold of $\theta = 0$), such that low-need customers would balk even when arriving at an empty system (L_1) and welfare would increase to $\Omega = 643$. However, rational customers anticipate this conflict of interest and disregard any information sent - leading to an “uninformative equilibrium” where welfare deteriorates to the *No-Information* level of $\Omega = 580$ (Proposition 4).

Table 1 Information Design and Welfare ($\Lambda = 4$, $p_H = 0.7$, $r_H = 300$, $r_L = 200$, $c = 50$, $v_H = v_L = 75$).

Arrival Seq.	No-Information				Full-Information				Partial-Information - Commitment				Partial-Information - Obedient			
	Queue	Utilities	Welfare		Queue	Balk	Utilities	Welfare	Queue	Balk	Utilities	Welfare	Queue	Balk	Utilities	Welfare
L_1, L_2, L_3, L_4	(L_1, L_2, L_3, L_4)	(150, 100, 50, 0)	300		(L_1, L_2)	L_3, L_4	(150, 100, 75, 75)	400	(L_1)	L_3, L_4, L_4	(150, 75, 75, 75)	375	$()$	L_3, L_3, L_3, L_4	75, 75, 75, 75	300
H_1, L_2, L_3, L_4	(H_1, L_2, L_3, L_4)	(250, 100, 50, 0)	400		(H_1, L_2)	L_3, L_4	(250, 100, 75, 75)	500	(H_1)	L_3, L_4	(250, 75, 75, 75)	475	(H_1)	L_3, L_3, L_4	(250, 75, 75, 75)	475
L_1, H_2, L_3, L_4	(L_1, H_2, L_3, L_4)	(150, 200, 50, 0)	400		(L_1, H_2)	L_3, L_4	(150, 200, 75, 75)	500	(L_1, H_2)	L_3, L_4	(150, 200, 75, 75)	500	(H_2)	L_3, L_3, L_4	(250, 75, 75, 75)	475
H_1, H_2, L_3, L_4	(H_1, H_2, L_3, L_4)	(250, 200, 50, 0)	500		(H_1, H_2)	L_3, L_4	(250, 200, 75, 75)	600	(H_1, H_2)	L_3, L_4	(250, 200, 75, 75)	600	(H_1, H_2)	L_3, L_4	(250, 200, 75, 75)	600
L_1, L_2, H_3, L_4	(L_1, L_2, H_3, L_4)	(150, 100, 150, 0)	400		(L_1, L_2, H_3)	L_4	(150, 100, 150, 75)	475	(L_1, H_3)	L_3, L_4	(150, 200, 75, 75)	500	(H_3)	L_3, L_3, L_4	(250, 75, 75, 75)	475
H_1, L_2, H_3, L_4	(H_1, L_2, H_3, L_4)	(250, 100, 150, 0)	500		(H_1, L_2, H_3)	L_4	(250, 100, 150, 75)	575	(H_1, H_3)	L_3, L_4	(250, 200, 75, 75)	600	(H_1, H_3)	L_3, L_4	(250, 200, 75, 75)	600
L_1, H_2, H_3, L_4	(L_1, H_2, H_3, L_4)	(150, 200, 150, 0)	500		(L_1, H_2, H_3)	L_4	(150, 200, 150, 75)	575	(L_1, H_2, H_3)	L_4	(150, 200, 150, 75)	575	(H_2, H_3)	L_3, L_4	(250, 200, 75, 75)	600
H_1, H_2, H_3, L_4	(H_1, H_2, H_3, L_4)	(250, 200, 150, 0)	600		(H_1, H_2, H_3)	L_4	(250, 200, 150, 75)	675	(H_1, H_2, H_3)	L_4	(250, 200, 150, 75)	675	(H_1, H_2, H_3)	L_4	(250, 200, 150, 75)	675
L_1, L_2, L_3, H_4	(L_1, L_2, L_3, H_4)	(150, 100, 50, 100)	400		(L_1, L_2, H_4)	L_3	(150, 100, 150, 75)	475	(L_1, H_4)	L_3, L_4	(150, 200, 75, 75)	500	(H_4)	L_3, L_3, L_4	(250, 75, 75, 75)	475
H_1, L_2, L_3, H_4	(H_1, L_2, L_3, H_4)	(250, 100, 50, 100)	500		(H_1, L_2, H_4)	L_3	(250, 100, 150, 75)	575	(H_1, H_4)	L_3, L_4	(250, 200, 75, 75)	600	(H_1, H_4)	L_3, L_4	(250, 200, 75, 75)	600
L_1, H_2, L_3, H_4	(L_1, H_2, L_3, H_4)	(150, 200, 50, 100)	500		(L_1, H_2, H_4)	L_3	(150, 200, 150, 75)	575	(L_1, H_2, H_4)	L_3	(150, 200, 150, 75)	575	(H_2, H_4)	L_3, L_4	(250, 200, 75, 75)	600
H_1, H_2, L_3, H_4	(H_1, H_2, L_3, H_4)	(250, 200, 50, 100)	600		(H_1, H_2, H_4)	L_3	(250, 200, 150, 75)	675	(H_1, H_2, H_4)	L_3	(250, 200, 150, 75)	675	(H_1, H_2, H_4)	L_3	(250, 200, 150, 75)	675
L_1, L_2, H_3, H_4	(L_1, L_2, H_3, H_4)	(150, 100, 150, 100)	500		(L_1, L_2, H_3, H_4)		(150, 100, 150, 100)	500	(L_1, H_3, H_4)	L_3	(150, 200, 150, 75)	575	(H_3, H_4)	L_3, L_4	(250, 200, 75, 75)	600
H_1, L_2, H_3, H_4	(H_1, L_2, H_3, H_4)	(250, 100, 150, 100)	600		(H_1, L_2, H_3, H_4)		(250, 100, 150, 100)	600	(H_1, H_3, H_4)	L_3	(250, 200, 150, 75)	675	(H_1, H_3, H_4)	L_3	(250, 200, 150, 75)	675
L_1, H_2, H_3, H_4	(L_1, H_2, H_3, H_4)	(150, 200, 150, 100)	600		(L_1, H_2, H_3, H_4)		(150, 200, 150, 100)	600	(L_1, H_2, H_3, H_4)		(150, 200, 150, 100)	600	(H_2, H_3, H_4)	L_3	(250, 200, 150, 75)	675
H_1, H_2, H_3, H_4	(H_1, H_2, H_3, H_4)	(250, 200, 150, 100)	700		(H_1, H_2, H_3, H_4)		(250, 200, 150, 100)	700	(H_1, H_2, H_3, H_4)		(250, 200, 150, 100)	700	(H_1, H_2, H_3, H_4)		(250, 200, 150, 100)	700
Expected Welfare	$\Omega = 580$				$\Omega = 620.5$				$\Omega = 634$				$\Omega = 643$			

4 Behavioral Experiments

We present the results from a series of experimental studies designed to test our main theoretical predictions under controlled laboratory conditions. Specifically, our model predicts that a service provider can improve social welfare with a properly designed partial information-sharing policy, and that welfare gains crucially depend on the provider’s ability to commit to a threshold policy.

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

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

Our theoretical framework offers sharp, testable predictions, but we note that departures from rationality may attenuate or shift these outcomes in practice. Importantly, the environment we study—built on key first principles that characterize service systems—creates a cognitively demanding setting for both providers and customers. This complexity can introduce decision noise, potentially dampening the predicted welfare gains from information provision (H1A: $\Omega^{Commit} > \Omega^{NoInfo}$). Furthermore, it may prevent customers from recognizing misaligned incentives with the provider, opening the door for information policies to generate positive welfare effects even in the absence of a commitment mechanism—contrary to theoretical predictions (H1B: $\Omega^{NoCommit} = \Omega^{NoInfo}$).

4.1 Implementation

4.1.1 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 Λ potential customers. At the beginning of each round, all customers are randomly and independently assigned a type: high-need with probability p_H , and low-need with $1 - p_H$. Customers arrive sequentially, according to randomly assigned and unknown indices $k \in \{1, \dots, \Lambda\}$, and decide whether to join the system or not based on available information. Customers who join the queue at position $q + 1$ earn $r_i - c(q + 1)$. Customers who do not join receive value of v_i 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. At the start of each round, each provider was randomly matched with Λ customers to mitigate reputation effects. Participants did not know their session size. The experimental treatments implemented in our study differ by the provider’s ability to send (credible) information, as discussed next.

4.1.2 Roadmap to Studies. We implement three treatments that vary the provider’s ability to share information with customers and to commit to an information-sharing policy. In the baseline *NoInfo* treatment, the provider cannot send any signals about the queue length. In *Commit*, the provider selects a threshold θ at the start of each round to generate binary signals $\varsigma = \{s, l\} = \{\text{short wait}, \text{long wait}\}$, sent to arriving customers. Customers receive s if the queue length is below θ upon their arrival, and l otherwise. The threshold θ is publicly announced and binding. In *NoCommit*, the provider privately selects a signal-generating threshold θ but publicly announces a potentially different threshold θ' . Customers do not observe θ , but know the provider may misreport. In contrast, in *Commit*, $\theta = \theta'$ by design. We compare these treatments with a *FullInfo* benchmark in §4.5. We present results from three studies. Study 1 implements all three treatments in a “small” population ($\Lambda = 4$), the minimal setting in which partial information outperforms full information (i.e., $\theta^* < q^*$) and the queue is symmetric around the Naor threshold (q^*), ensuring an unbiased prior by equally dividing favorable and unfavorable states. Study 2 re-implements *Commit* with an automated provider to control for suboptimal provider behavior observed in Study 1. Study 3 uses a larger population ($\Lambda = 8$), where the predicted welfare gap is stronger (providing stronger conditions for H1A), and where design features amplify the conflict of interest between providers and customers (providing stronger conditions for H1B).

4.1.3 Design Choices. We briefly discuss our main experimental design choices.

Human service providers. Our experiments feature both human customers and human service providers. We include a human provider even in the *NoInfo* treatment, where providers make no

decisions. This allows for a clean comparison across treatments, as it controls for the possibility that customer choices are (partially) driven by their social concern for the provider’s utility.

Strategy method. We use the strategy method to elicit customer choices (join or balk) before they learn their actual type for the round (see Appendix D). In the *Commit* and *NoCommit* treatments, we ask customers how they would respond to each possible signal, given communicated thresholds θ' . Widely used in experimental economics (e.g., Brandts 2011, Beer et al. 2022), the strategy method addresses two key challenges: First, it allows to fully understand customer behavior even in unlikely scenarios. Second, it maintains consistency with the theoretical assumption that customer rank k is unobservable. We need to ensure that participants cannot, even imperfectly, infer their rank k , e.g., from how long they have waited between rounds. The strategy method eliminates such inference. At the end of each round, the computer simulates the queue formation based on participant strategies and the realizations of random events (customer type and index); see Figure 2 for an example and Appendix D.1 for the rest treatments and Studies.

Feedback and learning. At the end of each round, all participants, regardless of their roles, receive detailed feedback, including thresholds, arrival order, types, signals, decisions, queue positions, and resulting utilities. While real-world service systems rarely offer such transparency, this design reinforces participant understanding of queue dynamics, signaling, and payoffs—giving theory its best shot. To aid learning, participants first complete practice rounds: 10 of *FullInfo* in Studies 1 and 2, and 5 rounds of *NoInfo* in Study 3. Additionally, participants completed comprehension quizzes where all questions had to be answered correctly to proceed with the experiment (see Appendix D.3).

Lying Detection Mechanism. The full feedback at the end of each round also means that customers in the *NoCommit* treatment can observe whether the provider’s implemented threshold θ matched the communicated threshold θ' . Although this may be unrealistic in practice, this gives customers the opportunity to detect the provider’s potential conflict of interest more easily, which gives theory its best shot.⁵

Parameters. In all studies, we choose parameter values $(r_H, r_L, c, p_H, \Lambda, v_H, v_L)$ such that all low-need customers theoretically join the system in *NoInfo* (see Proposition 1), thus creating a welfare loss that can be mitigated by a properly designed information policy (see §3.4).

4.1.4 Software, Recruitment, and Payment. We implemented the experiment in oTree (Chen et al. 2016), with 564 participants in total. Each subject took part in only one treatment. Participants were recruited from the subject pool of a large public university’s experimental lab. Upon arrival, they were randomly assigned to private cubicles and guided through on-screen instructions

⁵ One might argue that full feedback could act as an informal commitment device—discouraging providers from lying due to the risk of being detected. However, the extent of misreporting we observe suggests otherwise. In fact, lying is most pronounced in Study 3, precisely where the lie detection mechanism is strongest (see Appendix D).

Figure 2 Choice Elicitation and Queue Simulation: Experiment Interface for *Commit* in Study 1

Manager's selected strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait

Please select your strategy


	Long Wait	Short Wait
If you are H type (70% probability), you will	☒ Join ☐ Not Join	☒ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☒ Not Join	☒ Join ☐ Not Join

Your strategy


	☐	☐	☒	☐	☐
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait	

(a) Provider choice elicitation

(b) Focal customer choice elicitation



(c) Queue simulation: Focal customer arrives



(d) Queue simulation: End of round

Note: Following the strategy method, each round begins with service providers selecting their information policy (a), which is then communicated to customers. Customers respond by choosing their joining strategies (b). Once all strategies are set, the system randomly assigns customer indices and types, and executes decisions according to the elicited strategies. Participants then view simulations illustrating how arrivals, signals, and joining decisions unfold, visualized through animated graphics showing queue dynamics and resulting payoffs (c,d). See Appendix D.1 for the other treatments and Studies.

(see Appendix D.2). All participants could ask clarification questions before starting the main experiment. Each session lasted about 90 minutes. Participants received a fixed show-up fee—€5 in Studies 1 and 2, €7 in Study 3—and a bonus based on their average payoff over 40 rounds, converted at €0.10 per \$1 experimental dollar (Studies 1 and 2) or €0.50 per \$1 (Study 3). Total earnings ranged from €19–€23 (avg. €21.32) in Study 1, €19–€24 (avg. €21.53) in Study 2, and €7.44–€33.38 (avg. €19) in Study 3. All studies were pre-registered, including hypotheses, variables, analysis plans, and sample sizes. Pre-registrations are available at: <https://aspredicted.org/2bwq-73kp.pdf> and <https://aspredicted.org/x2dr-w8js.pdf> (Studies 1 and 2) and https://aspredicted.org/H92_D94 and https://aspredicted.org/MXN_MJZ (Study 3).

4.1.5 Data and Analysis. At the most granular level, for each round $t \in \{1, \dots, T\}$ and cohort $c \in \{1, \dots, C\}$, we observe the provider j 's implemented threshold θ_{jtc} in *Commit* and *NoCommit*, as well as the communicated threshold θ'_{jtc} in *NoCommit*. For each customer $k \in \{1, \dots, \Lambda\}$, we observe their strategy $A_{ktc} \equiv \{a_{ktc}^H, a_{ktc}^L\}$ in *NoInfo*, $A_{ktc} \equiv \{a_{ktc}^H(\varsigma = s|\theta_{jtc}), a_{ktc}^H(\varsigma = l|\theta_{jtc}), a_{ktc}^L(\varsigma = s|\theta_{jtc}), a_{ktc}^L(\varsigma = l|\theta_{jtc})\}$ in *Commit*, and $A_{ktc} \equiv \{a_{ktc}^H(\varsigma = s|\theta'_{jtc}), a_{ktc}^H(\varsigma = l|\theta'_{jtc}), a_{ktc}^L(\varsigma = s|\theta'_{jtc}), a_{ktc}^L(\varsigma = l|\theta'_{jtc})\}$ in *NoCommit*, where $a_{ktc}^i \in \{0, 1\}$ is such that 1 denotes *Join* and 0 denotes *Balk*. Rather than using realized welfare w_{ct} , we focus on the optimality of participant decisions by computing the *implied* welfare Ω_{ct} —the expected long-run welfare based on the observed joining strategies A_{ktc} and the implemented threshold θ_{jtc} . This measure removes randomness due to customer type and arrival order, allowing for cleaner comparisons with theoretical predictions. Appendix C details the computation of Ω_{ct} . For comparisons across treatments or with theory, we use session-level averages as the unit of analysis in our statistical tests (e.g., t-tests). We also report regression-based analyses, using standard errors clustered at the session level to account for dependencies in the data.

4.2 Study 1: Small Population ($\Lambda = 4$)

We implement three main treatments (*NoInfo*, *NoCommit*, *Commit*) in the task setting from §3.1, using the parameters from the numerical example in §3.6: $\Lambda = 4$, $p_H = 0.7$, $r_H = 300$, $r_L = 200$, $c = 50$, $v_H = v_L = 75$. We have a total of 230 participants, 8 sessions for *NoInfo*, 8 sessions for *Commit*, and 7 sessions for *NoCommit*. In each session, we have 2 cohorts of 5 (1 provider and 4 customers). Cohorts were randomly rematched in each round, and participants did not know their session size.

4.2.1 Choices: Customers. Across all treatments, and consistent with theory, we observe that high-need customers almost always join the queue (see Appendix E.1.1).

RESULT 1. High-need customers always join.

The joining behavior of low-need customers varies between treatments. It is indistinguishable from random joining in the *NoInfo* treatment (0.51 vs 0.5, $p = 0.795$), which highlights the complexity of making joining decisions without any information whatsoever. The joining rates of low-need customers are significantly less than theoretically predicted (0.512 vs 1, $p < 0.001$), yet they remain higher than socially optimal rates (0.51 vs 0, $p < 0.001$); see Proposition 1.

RESULT 2. Low-need customers join uniformly at random in *NoInfo*.

Our data further shows that the provision of signals in both *Commit* and *NoCommit* influence low-need customer behavior. In *Commit*, when receiving a short queue signal, customers are more likely to join than in *NoInfo* (0.914 vs. 0.512, $p < 0.001$), even though they join less than theoretically predicted (0.914 vs. 1, $p = 0.027$). Conversely, when receiving a long queue signal, customers are less likely to join than in *NoInfo* (0.100 vs. 0.512, $p < 0.001$), even though they join more than predicted

(0.100 vs. 0, $p = 0.007$). The effect of signals is strikingly similar in *NoCommit* - relative to the *NoInfo* baseline, customers are more likely to join after a short signal (0.891 vs. 0.512, $p < 0.001$), and less likely to join after a long signal (0.049 vs. 0.512, $p < 0.001$).⁶

RESULT 3. Signals are (equally) influential in *Commit* and *NoCommit*.

While signals are clearly influencing the joining behavior of low-need customers, the analysis thus far remains at an aggregate level and leaves unanswered the important question of whether signal strength depends on the signaling policy that providers communicate (i.e., threshold θ'). In Figures 3a and 3b, we present, for a given communicated threshold θ' and signal, the average customer joining probabilities. We also fit a simple logistic regression (see Appendix E.1.1, Table 4),

$$\mathbb{P}(\text{Join}_{it} = 1) = \Phi(\beta_0 + \beta_1 \cdot \text{NoCommit} + \beta_2 \cdot \theta' + \beta_3 \cdot \text{NoCommit} \cdot \theta' + \beta_4 \cdot \text{Round}), \quad (1)$$

where $\Phi(\cdot)$ is the Logistic function. Using *Commit* as the baseline, we include a dummy variable for *NoCommit*, the communicated threshold θ' , an interaction effect, and *Round* to capture possible learning effects. We estimate the model separately for short and long signals.

For short signals, we observe that low-need customers in *Commit* join less as θ' increases ($\hat{\beta}_2 = -1.54$, $p < 0.001$). Notably, we observe the same effect for *NoCommit* ($\hat{\beta}_3 = -0.22$, $p = 0.547$). For long signals, low-need customers in *Commit* join less for higher thresholds ($\hat{\beta}_2 = -1.02$, $p < 0.001$) and, again, we observe the same effect for *NoCommit* ($\hat{\beta}_3 = -0.14$, $p = 0.629$).


RESULT 4. Low-need customer joining rates in *Commit* and *NoCommit* decrease in θ' .

Overall, we find that customer behavior under *Commit* qualitatively aligns with theory. Surprisingly, *NoCommit* yields the same behavior as *Commit*. Regardless of the underlying mechanisms (which we discuss in §5.1), the customer choice patterns in both scenarios show the provider's ability to influence low-need customer behavior through strategic threshold selection, even without commitment. We found no evidence of customer learning in any treatment (see Table 4 in Appendix E.1.1).

4.2.2 Choices: Providers. In Figure 3c, we present the empirical distributions of thresholds communicated and implemented by the providers. In *Commit*, we observe that the average threshold is higher than theory predicts (2.26 vs. $\theta^* = 1$, $p = 0.002$), even in later rounds where we observe some modest learning (for $t = 21 - 40$: 2.15 vs. $\theta^* = 1$, $p = 0.004$). That is, providers send informative signals, but they send fewer long wait signals than optimal. Because customers tend to follow signals

⁶ Contrary to theory predictions, we observe that signals influence low-need customers in a similar fashion in *Commit* (0.891 vs. 1, $p = 0.041$) and *NoCommit* (0.049 vs. 1, $p < 0.001$), both for short wait signals (0.914 vs. 0.891, $p = 0.360$) and for long wait signals (0.100 vs. 0.049, $p = 0.109$).

Figure 3 Study 1: Decisions of Low-need Customers and Providers



(a) Low-need customers (short) (b) Low-need customers (long) (c) Providers: Threshold choices

even in the absence of commitment (Result 3), providers in *NoCommit* have an opportunity to influence customer joining behavior, contrary to what theory predicts. We observe that the average communicated threshold in *NoCommit* is indistinguishable from *Commit* (2.27 vs. 2.26, $p = 0.486$), and higher than optimal (2.27 vs. $\theta^* = 1$, $p < 0.001$). Similarly, the average implemented threshold in *NoCommit* is indistinguishable from *Commit* (2.03 vs. 2.26, $p = 0.250$), and is higher than what would be optimal (2.03 vs. $\theta^* = 1$, $p < 0.001$).

RESULT 5. Providers communicate and implement fewer long wait signals than optimal.

Intriguingly, Result 3 also presents an additional opportunity to steer customer behavior not present in *Commit*: the strategic communication and implementation of different thresholds. Indeed, 25.17% of the thresholds implemented by providers in *NoCommit* differ from those they communicated. The average implemented threshold is lower than communicated (2.03 vs. 2.27, $p = 0.077$), and the gap is even larger if we restrict the analysis to the second half of the data (1.85 vs. 2.21, $p = 0.023$).

RESULT 6. In *NoCommit*, providers communicate higher thresholds than implemented.


This suggests that providers behave strategically and in accordance with the conflict of interest in the model.

4.2.3 Social Welfare. Our results thus far show that signals are *influential* since customers react to them, and *informative* because they provide meaningful information about the queue length (even in *NoCommit*). Thus, our data provides evidence for the key mechanisms underlying the idea that wait signals can offer welfare gains. Indeed, Figure 4a shows that information design impacts welfare, though not in the direction predicted by theory. We observe that the implied social welfare

in *Commit* falls significantly short of theoretical predictions (613.5 vs. 634, $p < 0.001$), with no improvement over *NoInfo* (613.5 vs. 615.4, $p = 0.301$). In contrast, social welfare in *NoCommit* exceeds theoretical predictions (619.8 vs. 580, $p < 0.001$), and provides a modest improvement over *NoInfo* (619.8 vs. 615.4, $p = 0.111$) specifically in rounds $t = 21 - 40$, (621.73 vs. 615.78, $p = 0.032$), where providers learn to strategically communicate and implement different thresholds (Result 6).

RESULT 7. Signals in *Commit* do not improve welfare over *NoInfo*. Signals in *NoCommit* do.

Figure 4 Study 1: Decomposition of Social Welfare (incl. ● *CommitBOT*($\theta^* = 1$) from Study 2)



Why are there no welfare improvements in *Commit*, and modest welfare improvements in *NoCommit*? Several observations are notable here. First, welfare in *NoInfo* is substantially higher than predicted (615.43 vs 580, $p < 0.001$), leaving less room for improvement under our signaling treatments.⁷ Second, welfare gains in our signaling treatments are lower than what is possible (given the customer behavior that we observe), because of suboptimal provider behavior.

Figure 4b illustrates the welfare loss in the *NoCommit* treatment resulting from provider behavior. It breaks down threshold choices into three archetypes: *Honest*, *Overstate*, and *Understate*. When providers honestly communicate the implemented threshold (*Honest*: $\theta' = \theta$), they achieve the same welfare as in *NoInfo* (618.38 vs 615.43, $p = 0.191$). When providers communicate a higher threshold than they implement (*Overstate*: $\theta' > \theta$), effectively inflating the number of long-wait signals, they

⁷ While helpful for *NoInfo* (for which theory predicts 100% joining for the parameters implemented in our experiments), noisy customer behavior decreases welfare in *Commit*.

improve welfare over *NoInfo* (629.10 vs 615.43, $p = 0.007$) and approach the theoretical benchmark of the *Commit* treatment (629.10 vs 634.00, $p = 0.113$). Conversely, when providers communicate a lower threshold than they implement (*Understate*: $\theta' < \theta$), effectively inflating the number of short-wait signals, they achieve a lower welfare than in *NoInfo* (594.59 vs 615.43, $p = 0.001$). Overall, because providers overstate more often than they understate (Result 6), we observe higher welfare in *NoCommit* than in the other treatments (Result 7).

RESULT 8. In *NoCommit*, providers that communicate inflated thresholds achieve higher welfare.

Similarly, Figure 4c illustrates the welfare loss due to provider behavior for the *Commit* treatment, by decomposing the data “by threshold”. Unsurprisingly, non-informative thresholds fare poorly - when providers commit to always sending long wait signals regardless of the system state (i.e., $\theta' = 0$), they achieve the same welfare as in *NoInfo* (611.38 vs. 615.4, $p = 0.248$). When providers commit to always sending short wait signals (i.e., $\theta' = 4$), which induces more joining, they achieve even lower welfare than in *NoInfo* (603.83 vs. 615.4, $p = 0.013$). The high frequency of such non-informative thresholds in our data (Figure 3c) contributes to the low average welfare achieved in *Commit* (no higher than in *NoInfo*). Persuasive thresholds ($\theta' = \theta^* = 1$) provide welfare benefits: Figure 4c illustrates welfare gains over the *NoInfo* treatment in this case, although the increase is only modest (621.19 vs. 615.4, $p = 0.145$).

RESULT 9. In *Commit*, persuasive thresholds marginally increase welfare.

In general, implementing a good (i.e., *persuasive*) threshold in *Commit* suggests welfare improvements. However, this threshold is rarely picked by the providers, which points to a welfare improvement opportunity that we revisit in Study 2.

4.3 Study 2: Automated (and Persuasive) Providers

One of the most remarkable results of Study 1, is that signals can improve welfare when the provider has no commitment device (Result 7) - while *NoCommit* fares much better than theoretically predicted, *Commit* fares much worse. Customer behavior is unlikely to be the primary cause of lower welfare in *Commit* compared to *NoCommit*. After all, customer behavior is similar across both signal treatments (Results 3 and 4). Rather, Figure 4c suggests that poor provider choices are the main driver behind the low welfare in *Commit*. As a direct test of this conjecture, our second study implements a simple treatment, *CommitBOT*, which has $\Lambda = 4$ human customers arriving at a system with a “provider bot” that implements the theoretical “persuasive” threshold $\theta^* = 1$ throughout.

4.3.1 Design and Implementation. We followed the same design as in Study 1, with the only difference that the providers were automated. Overall, 64 subjects participated, organized in 8 sessions with 8 participants each.

4.3.2 Results. We first study whether automating the providers’s strategy in *CommitBOT* has a substantive impact on customer behavior. Consistent with Study 1 (Result 1), high-need customers join almost all the time (with 100% probability for “Short Wait” signals, and 99.7% for “Long Wait” signals). Similarly, when faced with a short wait signal, low-need customers in *CommitBOT* join the system with a high probability, comparable to customers facing the same $\theta = 1$ threshold in *Commit* (0.977 vs. 0.937, $p = 0.212$). Likewise, when given a “Long Wait” signal, low-need customers in *CommitBOT* join with a low probability, similar to those in *Commit* (0.171 vs. 0.220, $p = 0.216$).

RESULT 10. High and low-need customers joining behavior is similar in *Commit* and *CommitBOT*.

In terms of welfare, we observe that the consistent implementation of a persuasive threshold in *CommitBOT* yields higher welfare than *NoInfo* (Figure 4a: 626.26 vs. 615.43, $p = 0.003$), *Commit* (Figure 4a: 626.26 vs. 613.50, $p < 0.001$), and *NoCommit* (Figure 4a: 626.26 vs. 619.8, $p = 0.006$). Importantly, *CommitBot* does not improve on instances in *Commit* where providers do set the same persuasive threshold $\theta = 1$ (626.26 vs. 621.19, $p = 0.104$), nor does it improve over *NoCommit(Overstate)* (Figure 4c: 626.26 vs. 629.10, $p = 0.239$).

RESULT 11. The implementation of a persuasive threshold in *CommitBOT* increases welfare, but not more than the persuasive threshold in *Commit* and the strategically overstated thresholds in *NoCommit*.

Overall, we observe that *CommitBOT* improves welfare as it removes poorly set thresholds from the *Commit* treatment (Figure 4c). These results provide direct evidence that provider behavior is the main driver behind the poor welfare in *Commit*.

4.4 Study 3: Large Population ($\Lambda = 8$)

As real-world service systems typically involve the interactions of many customers, in this study we consider a larger customer population to examine the robustness of our results. Importantly, in this study, we make several design choices that highlight the conflict of interest between the provider and low-need customers, making it more salient to both parties. A more pronounced conflict environment strengthens the robustness of findings regarding the effects of signal treatments, where the provider may attempt to influence low-need customer behavior either through persuasion (*Commit*) or by potentially lying (*NoCommit*). We implement three main treatments (*NoInfo*, *NoCommit*, *Commit*) in the task setting from §3.1, using the following parameters: $\Lambda = 8$, $p_H = 0.65$, $r_H = r_L = 185$, $c = 40$, $v_L = 0$, and $v_H \ll 0$. As in the previous studies, parameter values are chosen such that all low-need customers theoretically join the system in *NoInfo* (see Proposition 1), thus creating a welfare loss that can be mitigated by a properly designed communication strategy (see §3.4).

We set the inside option utility equal across customer types ($r_H = r_L$), but distinguish between types through their outside option utilities. Specifically, low-need customers have a neutral outside

option ($v_L = 0$), while high-need customers face a strongly negative one ($v_H \ll 0$), such that low-need customers may opt out if the queue is long, whereas high-need customers always prefer joining the system. To reduce noise from customer decision-making, we automate decisions for high-need customers in line with theoretical predictions—specifically, that they always choose to join. This choice is reasonable in light of observed high-need behavior in our previous studies. Overall, this parametrization—where high-need customers always join but risk negative utility if they arrive late—makes the provider’s incentive to deter low-need customers from joining more salient to participants.⁸ Finally, while customers in Study 1 could only infer provider dishonesty in *NoCommit* with some effort, Study 3 implements a stronger lying-detection mechanism as we display the provider’s stated policy alongside their actual implemented policy (see Appendix D.1).

In total, 270 participants were included in our study and each subject participated in one treatment only. 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. For the following presentation of results, we preserve the numbering from Study 1 to allow for an easy and direct mapping between the results of both studies.

4.4.1 Choices: Customers. As in Study 1, the joining behavior of low-need customers is indistinguishable from random joining in the *NoInfo* treatment (0.571 vs 0.5, $p = 0.148$). Although low-need customers thus join significantly less than theoretically predicted *NoInfo* (0.571 vs 1, $p = 0.007$), they do join at higher than socially optimal rates (0.571 vs 0, $p = 0.004$).

RESULT 2’. Low-need customers join uniformly at random in *NoInfo*.

Do signals influence customer behavior? In *Commit*, low-need customers are more likely to join after a short signal than in *NoInfo* (0.880 vs. 0.571, $p < 0.001$), though less than predicted (0.880 vs. 1, $p = 0.003$). After a long signal, they are less likely to join than in *NoInfo* (0.238 vs. 0.571, $p = 0.001$), but more than predicted (0.238 vs. 0, $p = 0.002$). *NoCommit* shows a similar pattern: customers join more after short signals (0.837 vs. 0.571, $p < 0.001$) and less after long signals (0.308 vs. 0.571, $p = 0.001$). As in Study 1, signal responses are similar across *Commit* and *NoCommit*, for both short (0.880 vs. 0.837, $p = 0.093$) and long (0.238 vs. 0.308, $p = 0.099$) signals.

RESULT 3’. Signals are (equally) influential in *Commit* and *NoCommit*.

Does the information policy affect signal strength? Figures 5a and 5b show how customer joining probabilities vary with communicated threshold θ' and signal, and suggest that the information


⁸ Post-experimental survey results support the notion that study 3 presents a more salient conflict of interest (see Appendix F). Specifically, responses to item Q3, which measures perceived conflict of interest, show a clear and statistically significant increase across studies: participants in Study 1 reported the least conflict, followed by Study 2, with Study 3 showing the highest perceived conflict. This upward trend held consistently across customers, providers, and all treatment conditions (see Table 37 in Appendix F).

policy moderates the impact of signals on customer decisions. We formally test this using a logistic regression (Equation 1 from Study 1). We provide the full results in Appendix E.3.1 Table 24, and only briefly summarize the key findings here. As in Study 1, for short and long wait signals, the joining probability decreases in the communicated threshold in both *Commit* and *NoCommit* treatments. However, in contrast to Study 1, this probability decreases less sharply in the *NoCommit* treatment.

RESULT 4'. Communicated thresholds θ' decrease low-need customer joining, with a weaker effect in *NoCommit*.

A plausible result, given that in Study 3 the conflict of interest is more salient. Still, customers in *NoCommit* do respond to signals, suggesting providers could enhance welfare by overstating thresholds. We find no evidence of customer learning in the signaling treatments, and only a marginal decline in joining probabilities in *NoInfo* (see Table 24 in Appendix E.3.1).

Figure 5 Study 3: Decisions of Low-need Customers and Providers



(a) Low-need customers (short) (b) Low-need customers (long) (c) Providers: Threshold choices

4.4.2 Choices: Providers. Figure 5c presents the empirical distribution of thresholds communicated and implemented by the providers. In *Commit*, we observe that the average threshold is higher than predicted (2.98 vs. $\theta^* = 2$, $p = 0.008$), which persist with some modest learning (for $t = 21 - 40$: 2.83 vs. $\theta^* = 2$, $p = 0.004$). In other words, providers send informative signals, but fewer long wait signals than is optimal under a persuasive signaling policy. The average communicated threshold in *NoCommit* is indistinguishable from *Commit* (3.16 vs. 2.98, $p = 0.330$), and higher than optimal (3.16 vs. $\theta^* = 2$, $p = 0.004$). Similarly, the average implemented threshold in *NoCommit* is indistinguishable from *Commit* (2.46 vs. 2.98, $p = 0.180$), and is higher than what would be optimal (2.46 vs. $\theta^* = 2$, $p = 0.154$).

RESULT 5'. Providers communicate and implement fewer long wait signals than optimal.

As in Study 1, we observe that providers in *NoCommit* implement thresholds that are different from the communicated one 51.6% of the time. The average implemented threshold is lower than communicated (2.46 vs. 3.16, $p = 0.017$), and the gap is even larger if we restrict the analysis to the second half of the data (2.21 vs. 3.12, $p = 0.014$).

RESULT 6'. In *NoCommit*, providers communicate higher thresholds than implemented.

4.4.3 Social Welfare. Consistent with Study 1, we find that signals are influential and informative, thereby providing evidence for the key mechanisms that allow signals to generate welfare gains (even in *NoCommit*). Although welfare in *NoInfo* is substantially higher than theoretically predicted (172.83 vs 40, $p = 0.005$), Figure 6a shows that information design improves welfare. We observe that implied social welfare in *Commit* falls significantly short of theoretical predictions (200.70 vs. 250, $p < 0.001$), yet provides improvement over *NoInfo* (200.70 vs. 172.83, $p = 0.026$), which is sustained when looking only at the second half of the data, from $t = 21 - 40$, (204.53 vs. 181.52, $p = 0.062$). Similarly, social welfare in *NoCommit* exceeds theoretical predictions (198.70 vs. 40, $p < 0.001$), and provides an improvement over *NoInfo* (198.70 vs. 172.83, $p = 0.015$), which is sustained when looking only at the second half of the data, from $t = 21 - 40$, (202.55 vs. 181.52, $p = 0.017$).

RESULT 7'. Signals in *Commit* and *NoCommit* improve welfare.

Figure 6 Study 3: Decomposition of Social Welfare




Figure 6b decomposes providers' threshold choices. When providers honestly communicate the implemented threshold (*Honest*: $\theta' = \theta$), they achieve higher welfare than in *NoInfo* (195.08 vs 172.83, $p = 0.022$). Similarly, when providers communicate a higher threshold than they implement (*Overstate*: $\theta' > \theta$), i.e., inflate the number of long-wait signals, they improve welfare over *NoInfo* (218.91 vs 172.83, $p < 0.001$) and approach the theoretical benchmark of the *Commit* treatment (218.91 vs 250, $p < 0.001$). Conversely, when providers communicate a lower threshold than they implement (*Understate*: $\theta' < \theta$), i.e., inflate the number of short-wait signals, they achieve the same welfare as in *NoInfo* (169.06 vs 172.83, $p = 0.383$).

RESULT 8'. In *NoCommit*, providers that communicate inflated thresholds achieve higher welfare.

In a similar spirit, Figure 6c visualizes welfare loss from provider behavior for the *Commit* treatment, by decomposing the data “by threshold”. Not surprisingly, non-informative thresholds fare poorly - when providers commit to always sending long wait signals regardless of the system state (i.e., $\theta' = 0$), they achieve the same welfare as in *NoInfo* (196.46 vs. 172.83, $p = 0.091$). Similarly, when providers commit to always sending short wait signals (i.e., $\theta' = 8$), which induces more joining, they achieve the same welfare as in *NoInfo* (192.26 vs. 172.83, $p = 0.096$). It stands to reason that persuasive thresholds (i.e., $\theta' = \theta^* = 2$) provide welfare benefits. Indeed, the persuasive threshold increases welfare compared to the *NoInfo* treatment (207.66 vs. 172.83, $p = 0.012$).

RESULT 9'. In *Commit*, providers that implement a persuasive threshold improve welfare.

4.5 Full-Information Benchmark: The Observable Queue

An important theoretical prediction is that neither extreme—*NoInfo* nor *FullInfo*—is optimal. Instead, the optimal signaling strategy involves conveying partial information (see §3.5). To study how binary signals compare to a full information benchmark, we first look at the 10 practice rounds that participants in Studies 1 and 2 completed under a *FullInfo* treatment. This provided a simple reference point for assessing performance under fully observable queues. Social welfare in these practice rounds consistently fell short of theoretical predictions, and all main treatments in both studies outperformed their respective *FullInfo* practice counterpart (see Appendices E.1.4, E.2.3).

Because our *FullInfo* data is from unincentivized practice rounds intended primarily for learning, we also use the theoretical *FullInfo* benchmark (see Proposition 2) as a more conservative point of comparison. We find that binary signaling achieves or exceeds the theoretical welfare predicted under a *FullInfo* environment across multiple treatments and studies. In Study 1, both *Commit*($\theta^* = 1$) and *NoCommit* yield outcomes consistent with the *FullInfo* benchmark (621.19 vs. 620.5, $p = 0.441$ and 619.83 vs. 620.5, $p = 0.327$, respectively), while *NoCommit*(*Overstate*) exceeds it significantly (629.10 vs. 620.5, $p = 0.034$). In Study 2, *CommitBOT* also surpasses the *FullInfo* benchmark (626.26 vs.

620.5, $p = 0.005$). Similarly, in Study 3, both *Commit* and *NoCommit* match the *FullInfo* level (200.7 vs. 199.6, $p = 0.416$ and 198.7 vs. 199.6, $p = 0.412$), while *Commit*($\theta^* = 2$) and *NoCommit*(*Overstate*) outperform it (207.66 vs. 199.6, $p = 0.095$ and 218.91 vs. 199.6, $p = 0.001$, respectively).

5 Behavioral Mechanisms and Managerial Implications


Overall, we find that both customers and providers deviate from the perfect rationality assumption, yet they do exhibit some strategic reasoning. Customers behave randomly in the absence of queue information in *NoInfo*, and appear insensitive to conflicts of interest in *NoCommit*. Still, they respond to signals in ways that qualitatively align with theoretical predictions—adjusting their behavior based on how signals map to queue length. Providers, while rarely adopting fully persuasive signaling policies in *Commit*, often choose reasonably informative ones and frequently engage in strategic misreporting when given the opportunity in *NoCommit*. We find that information provision is beneficial for welfare, albeit to a lesser extent than predicted by theory. Importantly, even non-binding communication leads to improvements.

5.1 Behavioral Mechanisms

While our experimental design does not aim to isolate behavioral mechanisms, observed patterns align with established findings in behavioral economics and operations. In particular, the aggregate behavior of customers and providers fits well within the Quantal Response Equilibrium (QRE) framework (Goeree et al. 2016), consistent with evidence of bounded rationality in strategic communication games and its growing use in behavioral operations (see §2). QRE relaxes the assumption of perfect optimization by allowing for the possibility of mistakes: better strategies are more likely to be chosen, but not always. This is modeled via a mean-zero noise term in expected utility, capturing systematic decision errors while preserving responsiveness to information. For instance, assuming independent noise terms follow a mean-zero Gumbel distribution with scale parameter $\beta > 0$ yields a logit form for choice probabilities (McFadden 1981). Using the *NoInfo* setting as an example, the joining probability for low-need customers is then given by:

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

The parameter β reflects the degree of *bounded rationality* (Chen et al. 2012, Huang et al. 2013). As $\beta \rightarrow 0$, choices converge to full rationality; as $\beta \rightarrow \infty$, choices become uniformly random. We estimate β with our experimental data, providing a quantitative measure of bounded rationality. Appendix G provides details for all treatments and structural estimates from our data. To illustrate, Figure 7 shows predicted choice probabilities under the estimated QRE model for all three treatments of Study 3; results for all studies are reported in Appendix G.

Figure 7 QRE Predictions: Low-need Customers and Providers (Full Models from Table 39 in Appendix G)

(a) Low-need customers (short) (b) Low-need customers (long) (c) Providers: Threshold choices

Originally proposed by Huang et al. (2013) for queueing settings, QRE accounts for customers' difficulty in estimating queue length from available information and providers' imperfect welfare optimization. Our results are consistent with this view. For instance, low-need customers under *NoInfo* appear to behave randomly (Result 2'), reflecting difficulty in forming queue length expectations in the absence of information. In *Commit*, binary signals reduce this burden, as indicated by low-need customers joining more often after short signals and less after long ones (Figures 7a and 7b). Their joining probability also decreases with higher communicated thresholds θ' (Result 4')—a pattern QRE predicts, as errors become less likely when utility differences between choices grow (Goeree et al. 2016). Provider behavior under *Commit* also matches QRE predictions: the distribution of thresholds aligns with welfare-maximizing ones based on customers' responses (Figure 7c).

Under *NoCommit*, while signals are less informative, they still guide behavior (Result 3')—Figures 7a and 7b show how customers respond differently to short vs. long signals, though less strongly than under *Commit* (Result 4'). QRE explains this muted responsiveness as stemming from customers' uncertainty about the true threshold θ , leading to flatter expected payoffs. Providers also tend to inflate θ' but not maximally (Result 6'). We find that introducing a lying cost in the QRE model captures this behavior, which is consistent with truth-telling preferences documented in related literature (Rosenbaum et al. 2014, Abeler et al. 2019, Özer et al. 2011, Rodriguez et al. 2025).

QRE also helps explain the welfare deviations from theoretical predictions observed in our data. In *NoInfo*, random errors reduce overjoining, boosting welfare—supporting Huang et al. (2013)'s theoretical prediction that bounded rationality can mitigate congestion. This effect is particularly pronounced because estimating queue lengths is most cognitively demanding in *NoInfo*. In *Commit*, providers often fail to choose persuasive thresholds, and even when they do (e.g., *CommitBOT*),

boundedly rational customers are imperfectly persuaded. Interestingly, this means that customer perfect rationality would enhance persuasion and increase welfare under commitment. In *NoCommit*, lying aversion and bounded rationality yield informative communication with welfare outcomes close to those in *Commit*. Overall, these behavioral mechanisms significantly shape how information design impacts welfare in queues. Additional details and discussions are presented in Appendix G.

5.2 Managerial Implications

Our results offer insights for information design in service systems. At a high level, our results indicate that firms can use information sharing as an effective managerial lever to mitigate customer overjoining in service systems. Our results suggest that firms can treat formal commitment mechanisms and behavioral mechanisms as substitutes when designing information policies. In settings with fully rational customers, commitment devices are crucial—shared information alone does not influence behavior when provider and customer incentives are misaligned, as rational customers disregard untrustworthy signals. However, with boundedly rational customers, the effectiveness of information policies depends less on formal commitment and more on whether customers can perceive and respond to conflicts of interest, as well as the extent to which providers share information truthfully. This behavioral perspective broadens the range of effective policy tools available to firms.

A natural implication of the above is that firms do not necessarily need to invest in costly commitment devices—such as audit trails or blockchain technology—for information design interventions to work. Indeed, our findings show that a lack of commitment does not necessarily disrupt communication; rather, it gives providers the option to inflate delay announcements as an additional lever to mitigate overjoining behavior. This also falls in line with the conventional wisdom of “under-promising and over-delivering”, which can enhance customer satisfaction by creating a positive surprise for those who choose to join the system (Yu and Smith 2020). While we do not endorse deception—given the ethical, legal, and reputational implications—our results may help explain potential practices in the field (Lazarus 2020, Fisher 2025). Clearly, more research on this delicate subject is in order.

Finally, our results show that when commitment mechanisms are in place, firms should consider using binary signals that convey vague information rather than fully informative ones. Binary signals are not only easier to implement, but also serve as effective persuasive tools. While boundedly rational customers may struggle to interpret vague information—resulting in muted behavioral responses—we find that such signals can still match or even exceed the welfare outcomes of full-information policies. This suggests that such strategic obfuscation can be both practical and robust in the presence of behavioral noise.

6 Conclusions

Information design, as a managerial lever to mitigate customer overuse and improve welfare in service systems, is receiving increasing attention from service operations scholars and practitioners alike. Our study offers an initial empirical test of key assumptions in the theoretical information design literature and anecdotal evidence from practice. We show that carefully crafted binary signals can shape customer behavior and enhance social welfare—even without formal commitment mechanisms.

Our design choices have limitations that open avenues for future research. We focused on a parameter space where low-need customers in our experiments behave randomly in the absence of information. Future work could explore alternative parameter settings to test whether such randomness persists and examine the conditions under which strategically withholding information—thereby inducing random joining behavior—can achieve welfare outcomes comparable to those of more complex signaling policies. In our experiment, we used qualitative signals (e.g., “long wait”); future studies could investigate how different forms of signal framing—such as queue lengths, estimated wait times, or action-oriented messages—affect the quality of customer decision-making. Action-recommendation messages (e.g., “check in online to reduce your wait”) may prove more effective by translating system state into concrete, relevant choices. Our model focuses on parameter settings where high-need customers always join; future work could relax this to examine whether bounded rationality unintentionally deters high-need customers, highlighting the importance of robust targeting information policies. Finally, we study information design in a simple first-come, first-served queue, future research could study prioritization settings to examine how information sharing complements such policies.

Finally, we hope to inspire new theoretical work on information design in queueing systems grounded on bounded rationality. Additionally, given the challenges of replicating queueing environments in laboratory settings, we hope our study helps advance more experimental research on queueing in operations.

References

- Abeler, J., D. Nosenzo, and C. Raymond (2019). Preferences for truth-telling. *Econometrica* 87(4), 1115–1153.
- Akşın, 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*.
- 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. Irvani, 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*.
- Argon, N. T. and S. Ziya (2009). Priority assignment under imperfect information on customer type identities. *Manufacturing & Service Operations Management* 11(4), 674–693.
- Au, P. H., O. Kwon, and K. K. Li (2023). A simple experiment on simple bayesian persuasion. *Working Paper*.
- 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.
- Blume, A., E. K. Lai, and W. Lim (2020). Strategic information transmission: A survey of experiments and theoretical foundations. *Handbook of experimental game theory*, 311–347.
- Bramwell, K. (2023). A&e crisis: Staff 'broken' as 2023 looks set to be worst year ever. <https://www.bbc.com/news/health-64136691>. Accessed: 2024-08-01.
- Brandts, J., C.-G. (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics* 14, 375–398.
- Bricker, T. (2024). Actual versus posted wait times at walt disney world. <https://www.disneytouristblog.com/actual-versus-posted-wait-times-disney-world/>. Accessed: 2025-07-13.
- 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.
- Cai, H. and J. T.-Y. Wang (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior* 56(1), 7–36.
- 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. *Queueing 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.
- Coricelli, G., A. Aristidou, and A. Vostroknutov (2023). Incentives or persuasion? an experimental investigation. *Available at SSRN*.
- 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.
- Drakopoulos, K., I. Lo, and J. Mulvany (2022). Blockchain mediated persuasion. *USC Marshall School of Business Research Paper Sponsored by iORB*.
- 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 queueing processes. *Econometrica: Journal of the Econometric Society*, 81–92.
- EU-Commission (2020). European pillar of social rights. <https://op.europa.eu/en/publication-detail/-/publication/9b16d965-090e-11eb-bc07-01aa75ed71a1>. Accessed: 2025-07-13.
- Fisher, R. (2025). Are contact centers really “experiencing unusually high call volumes”? martin lewis is on the case. <https://www.cxtoday.com/contact-center/are-contact-centers-really-experiencing-unusually-high-call-volumes-martin-lewis-is-on-the-case/>. Accessed: 2025-07-13.

- Frazelle, A. E. and E. Katok (2024). Paid priority in service systems: Theory and experiments. *Manufacturing & Service Operations Management*.
- 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.
- GovUk (2024). Shrewd waitless. <https://www.applytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/537879466654210>. Accessed: 2024-06-17.
- 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.
- Hassin, R. (1995). Decentralized regulation of a queue. *Management Science* 41(1), 163–173.
- Hassin, R. (2016). *Rational queueing*. CRC press.
- Hassin, R. and M. Haviv (2003). *To queue or not to queue: Equilibrium behavior in queueing systems*, Volume 59. Springer Science & Business Media.
- Haviv, M. (2014). Regulating an m/g/1 queue when customers know their demand. *Performance Evaluation* 77, 57–71.
- 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.
- Hyde, J. (2018). Supreme court finds hospital negligent over misleading waiting time. <https://www.lawgazette.co.uk/law/supreme-court-finds-hospital-negligent-over-misleading-waiting-time/5067884.article>. Accessed: 13-Feb-2025.
- 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.
- Kawagoe, T. and H. Takizawa (2009). Equilibrium refinement vs. level-k analysis: An experimental study of cheap-talk games with private information. *Games and Economic Behavior* 66(1), 238–255.

- Kremer, M. and L. Debo (2016). Inferring quality from wait time. *Management Science* 62(10), 3023–3038.
- Kumar, P., M. U. Kalwani, and M. Dada (1997). The impact of waiting time guarantees on customers’ waiting experiences. *Marketing science* 16(4), 295–314.
- Lazarus, D. (2020). “unusually high call volume”? you’re being lied to, say experts. <https://www.latimes.com/business/story/2020-11-10/column-coronavirus-high-call-volume>. Accessed: 2025-07-13.
- Lin, C.-A., K. Shang, and P. Sun (2023). Wait time-based pricing for queues with customer-chosen service times. *Management Science* 69(4), 2127–2146.
- Lingenbrink, D. and K. Iyer (2019). Optimal signaling mechanisms in unobservable queues. *Operations research* 67(5), 1397–1416.
- Lusa (2023). Demand for hospital emergency room peaking. <https://www.theportugalnews.com/news/2023-05-24/demand-for-hospital-emergency-room-peaking/77923>. Accessed: 2024-08-01.
- McFadden, D. (1981). Econometric models of probabilistic choice. *Structural analysis of discrete data with econometric applications* 198272.
- Mendelson, H. and S. Whang (1990). Optimal incentive-compatible priority pricing for the m/m/1 queue. *Operations research* 38(5), 870–883.
- 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.
- Narayanan, R. (2023). Enhancing transparency in patient wait time: Exploring blockchain solutions in health-care queue tracking innovation. <https://www.medigy.com/news/blogs/enhancing-transparency-in-patient-wait-time-exploring-blockchain-solutions-in-healthcare-queue-tracking-innovation/>. Accessed: 2025-07-13.
- Nguyen, Q. (2017). Bayesian persuasion: evidence from the laboratory. *Work. Pap., Utah State Univ., Logan*.
- NHS (2023). The nhs choice framework: What choices are available to me in the nhs? <https://www.gov.uk/government/publications/the-nhs-choice-framework/the-nhs-choice-framework-what-choices-are-available-to-me-in-the-nhs>. Accessed: 2024-06-03.
- Özer, Ö., Y. Zheng, and K.-Y. Chen (2011). Trust in forecast information sharing. *Management Science* 57(6), 1111–1137.
- Rodriguez, E. A., R. Ibrahim, and D. Zhan (2025). On customer (dis-) honesty in unobservable queues: The role of lying aversion. *Management Science* 71(1), 844–860.
- 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.

- Scheele, L. M., U. W. Thonemann, and M. Slikker (2018). Designing incentive systems for truthful forecast information sharing within a firm. *Management Science* 64(8), 3690–3713.
- 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.
- Singh, S. P., M. Delasay, and A. Scheller-Wolf (2023). Real-time delay announcement under competition. *Production and Operations Management* 32(3), 863–881.
- Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management* 10(4), 566–589.
- Wang, J. T.-y., M. Spezio, and C. F. Camerer (2010). Pinocchio’s pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American economic review* 100(3), 984–1007.
- Webb, E., Q. Yu, and K. Bretthauer (2020). Linking delay announcements, abandonment, and service time. *Working Paper*.
- Wu, W. and B. Ye (2023). Competition in persuasion: An experiment. *Games and Economic Behavior* 138, 72–89.
- 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. and J. Smith (2020). When providing wait times, it pays to underpromise and overdeliver. *Harvard Business Review*. Accessed: 2025-07-13.
- 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.
- Özer, , N. Subramanian, and Y. Wang (2018). Information sharing, advice provision, or delegation: What leads to higher trust and trustworthiness? *Management Science* 64(1), 474–493.

Appendices

A Technical Proofs

A.1 Proposition 1 (No-Information)

PROOF.

In this proof, we begin by analyzing the self-interested joining behavior of customers. Then, we study the socially optimal (i.e., social welfare maximizing) behavior, and conclude with a comparison of the two.

(Self-Interested Joining Behavior). Consider a tagged low-need customer that joins with probability α' , when the other low-need customers join with probability α . Since all high-need 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 $w \in \{0, 1, \dots, \Lambda - 1\}$. 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 $\Lambda - 1, \Lambda - 2, \dots, w + 1$; 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 $\mathbb{E}[U(\alpha', \alpha)] = (1 - \alpha')v_L + \alpha'(r_L - c(\mathbb{E}[Q] + 1))$. To find her 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 $r_L - v_L - c(\mathbb{E}[Q] + 1)$. Let $r'_L \triangleq r_L - v_L$, and the root of $r'_L - c(\mathbb{E}[Q] + 1) = 0$ be

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

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 $\frac{(\Lambda-1)p_H}{2} + 1 \geq \frac{r'_L}{c}$.
- 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 $\frac{\Lambda-1}{2} + 1 \leq \frac{r'_L}{c}$.
- 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 $\frac{(\Lambda-1)p_H}{2} + 1 < \frac{r'_L}{c} < \frac{(\Lambda-1)}{2} + 1$.

(Socially Optimal Joining Behavior). Let $\alpha_{soc} \in [0, 1]$ denote a fixed joining probability for low-need customers. Based on this, we can compute the expected social welfare as follows:

$$\begin{aligned}\Omega &= \mathbb{E} \left[\sum_{k=1}^{\Lambda} U_k \right] = r_H \mathbb{E}[J_H] + r_L \mathbb{E}[J_L] - c \mathbb{E} \left[\sum_{k=1}^J k \right] + v_L (\Lambda - \mathbb{E}[J]) \\ &= r_H \mathbb{E}[J_H] + r_L (\mathbb{E}[J] - \mathbb{E}[J_H]) - \frac{c}{2} (\mathbb{E}[J^2] + \mathbb{E}[J]) + v_L (\Lambda - \mathbb{E}[J]),\end{aligned}$$

where $J = J_H + J_L$ is a random variable that represents the total number of customers that *join* the system. Similarly, J_H (respectively, J_L) are random variables that represent the total number of high-need (respectively, low-need) customers that join the system. In particular, since all high-need customers join, we have that $J \sim \mathcal{B}(\Lambda, p_H + (1-p_H)\alpha_{soc})$ is a binomial random variable. Moreover, notice that $J_H | J \sim \mathcal{B}(J, \frac{p_H}{p_H + (1-p_H)\alpha_{soc}})$, such that $J_H \sim \mathcal{B}(\Lambda, p_H)$ from the property of conditional binomials. From this, we have that $\mathbb{E}[J] = \Lambda(p_H + (1-p_H)\alpha_{soc})$, $\mathbb{E}[J_H] = \Lambda p_H$, and $\mathbb{E}[J^2] = \Lambda(p_H + (1-p_H)\alpha_{soc}) + \Lambda(\Lambda-1)(p_H + (1-p_H)\alpha_{soc})^2$. Based on this, after some algebra, we have that:

$$\begin{aligned}\frac{\partial \Omega}{\partial \alpha_{soc}} &= \Omega'(\alpha_{soc}) = \Lambda(1-p_H)(r'_L - c(1 + (\Lambda-1)(p_H + (1-p_H)\alpha_{soc}))), \\ \Omega'(0) &= \Lambda(1-p_H)(r'_L - c(1 + p_H(\Lambda-1))), \\ \Omega'(1) &= \Lambda(1-p_H)(r'_L - c\Lambda) < 0, \\ \frac{\partial^2 \Omega}{\partial \alpha_{soc}^2} &= \Omega''(\alpha_{soc}) = -c(\Lambda-1)\Lambda(1-p_H)^2 < 0.\end{aligned}$$

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

(Over-Joining). From the above results, we first note that since $\frac{(\Lambda-1)p_H}{2} + 1 < 1 + p_H(\Lambda-1)$, if $\alpha^* = 0$ then $\alpha_{soc}^* = 0$. Now, consider the joining probability $\alpha_{soc}^* = \frac{r'_L - c(1 + p_H(\Lambda-1))}{c(\Lambda-1)(1-p_H)} < 1$, and the probability $\alpha^* = \frac{2r'_L - c(2 + p_H(\Lambda-1))}{c(\Lambda-1)(1-p_H)}$. After some simple algebra, it is easy to see that $\alpha^* > \alpha_{soc}^* \iff r'_L > c$, which is our assumption by construction. This also means that for $\alpha^* = 1$ we have that $\alpha^* > \alpha_{soc}^*$. ■

A.2 Proposition 2 (Full-Information)

PROOF.

In this proof, we begin by analyzing the self-interested joining behavior of customers. Then, we study the socially optimal (i.e., social welfare maximizing) behavior, and conclude with a comparison of the two.

(Self-Interested Joining Behavior). Consider a tagged low-need customer that joins with probability $\alpha'_q = (\alpha'_0, \alpha'_1, \dots, \alpha'_{\Lambda-1})$, when high-need customers always join and the other low-need customers follow the strategy $\alpha_q = (\alpha_0, \alpha_1, \dots, \alpha_{\Lambda-1})$. When the tagged customer arrives, she knows the number of people in the system q , such that her expected utility is $\mathbb{E}[U(\alpha'_q, \alpha_q)|q] = (1 - \alpha'_q)v_L + \alpha'_q(r_L - c(q+1))$. To find her best response against α_q , the tagged customer has to find the α'_q that maximizes $\mathbb{E}[U(\alpha'_q, \alpha_q)|q]$. Note that the function $\mathbb{E}[U(\alpha'_q, \alpha_q)|q]$ is linear with respect to α'_q , so the tagged customer bases her decision on the sign of the quantity $r_L - v_L - c(q+1)$. Recall that $r'_L = r_L - v_L$. Based on this, the set of best responses against α_q , $BR(\alpha_q|q) = \arg \max_{\alpha'_q} \mathbb{E}[U(\alpha'_q, \alpha_q)|q]$, is given by

$$BR(\alpha_q|q) = \begin{cases} 0 & \text{if } \frac{r'_L}{c} < q+1 \\ [0, 1] & \text{if } \frac{r'_L}{c} = q+1 \\ 1 & \text{if } \frac{r'_L}{c} > q+1. \end{cases}$$

Note that $BR(\alpha_q|q)$ does not depend on α_q . It follows that low-need customers will join whenever their position in the queue (i.e., $q+1$) does not exceed $q^* = \lfloor \frac{r'_L}{c} \rfloor$.

(Socially Optimal Joining Behavior). Let $q_{soc} \in \{0, 1, \dots, \Lambda\}$ denote a fixed threshold, such that low-need customers join with probability 1 if their position in queue is $q+1 \leq q_{soc}$, and balk with probability 1 if $q+1 > q_{soc}$. Based on this, we can compute the expected social welfare as follows:

$$\begin{aligned} \Omega(q_{soc}) &= r_H(q_{soc}p_H + \mathbb{E}[N_H]) + r_Lq_{soc}(1 - p_H) \\ &\quad - \frac{c}{2}(\mathbb{E}[J(q_{soc})^2] + \mathbb{E}[J(q_{soc})]) + v_L(\Lambda - \mathbb{E}[J(q_{soc})]) \\ &= r_H(q_{soc}p_H + \mathbb{E}[N_H]) + r_Lq_{soc}(1 - p_H) \\ &\quad - \frac{c}{2}(q_{soc}^2 + 2q_{soc}\mathbb{E}[N_H] + \mathbb{E}[N_H^2] + q_{soc} + \mathbb{E}[N_H]) + v_L(\Lambda - \mathbb{E}[J(q_{soc})]). \end{aligned}$$

Since in this case the first q_{soc} customers join, and after the threshold q_{soc} only high-need customers join, we have that $J(q_{soc}) = q_{soc} + N_H$ is a random variable that represents the total number of customers that join according to strategy q_{soc} , and $N_H \sim \mathcal{B}(\Lambda - q_{soc}, p_H)$ is a binomial random variable that represents the total number of high-need customers that join after the threshold q_{soc} . From this, we have that $\mathbb{E}[N_H] = (\Lambda - q_{soc})p_H$, and $\mathbb{E}[N_H^2] = (\Lambda - q_{soc})p_H + (\Lambda - q_{soc})(\Lambda - q_{soc} - 1)p_H^2$. Now, consider the following expression:

$$\Omega(q_{soc}) - \Omega(q_{soc} - 1) = (1 - p_H)(r'_L - c(\Lambda p_H + (1 - p_H)q_{soc})).$$

Note that for $q_{soc} = 0$ we have that $r'_L - c\Lambda p_H \geq 0 \iff p_H \leq \frac{r'_L}{c\Lambda}$, and for $q_{soc} = \Lambda$ it becomes $r'_L - c > 0$, which is always the case. Moreover, notice that $r'_L - c(\Lambda p_H + (1 - p_H)q_{soc})$ decreases strictly in q_{soc} . It follows that whenever $p_H \leq \frac{r'_L}{c\Lambda}$ there is a unique threshold q_{soc}^* such that $q_{soc} \leq \frac{r'_L - c\Lambda p_H}{c(1 - p_H)}$ for $q_{soc} \leq q_{soc}^*$, and $q_{soc} > \frac{r'_L - c\Lambda p_H}{c(1 - p_H)}$ for $q_{soc} > q_{soc}^*$. It follows that $\Omega(q_{soc}) - \Omega(q_{soc} - 1) \geq 0$ for $q_{soc} \leq q_{soc}^*$, and $\Omega(q_{soc}) - \Omega(q_{soc} - 1) < 0$ for

$q_{soc} > q_{soc}^*$. We conclude that whenever $p_H \leq \frac{r'_L}{c\Lambda}$, $\Omega(q_{soc})$ is unimodal with a maximum at $q_{soc}^* = \lfloor \frac{r'_L - c\Lambda p_H}{c(1-p_H)} \rfloor$. Finally, whenever $p_H > \frac{r'_L}{c\Lambda}$ we have that $\Omega(q_{soc}) - \Omega(q_{soc} - 1) < 0$ for all q_{soc} , and thus $q_{soc}^* = 0$.

(Over-Joining). From the above results, whenever $p_H = 0$ we have that $q_{soc}^* = \lfloor \frac{r'_L}{c} \rfloor = q^*$. Moreover, whenever $p_H > 0$, for $p_H \leq \frac{r'_L}{c\Lambda}$, we have that $\frac{r'_L - c\Lambda p_H}{c(1-p_H)} < \frac{r'_L}{c} \iff r'_L - c\Lambda < 0$, and for $p_H > \frac{r'_L}{c\Lambda}$ we have $0 < \frac{r'_L}{c}$. That is, overall we have that $q_{soc}^* \leq q^*$. ■

A.3 Lemma 1

PROOF. Let $\theta \in \{0, \dots, \Lambda\}$ be a fixed threshold. Consider a tagged low-need customer that joins with probabilities α'_s and α'_l , when high-need customers always join and the other low-need customers join with probabilities α_s and α_l . Given the parameter space in the game, $\frac{r'_L}{c} \geq \frac{\Lambda-1}{2} + 1$, it follows that $r'_L - c(\mathbb{E}[Q] + 1) \geq 0$ 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 $\mathbb{E}[Q|\zeta = s] \leq \mathbb{E}[Q]$ for any threshold θ , it follows that $r'_L - c(\mathbb{E}[Q|\zeta = s] + 1) \geq 0$ 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 short signal, irrespective of how others behave. It follows that $\alpha_s^* = 1$.

Now, when low-need customers receive a long signal, they know that they are in a system with $q \in \{\theta, \theta + 1, \dots, \Lambda - 1\}$ customers. This is because low-need customers that receive a short signal join the system with probability $\alpha_s^* = 1$ in equilibrium. Based on this, and since high-need customers 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 $\theta + j$ ($j \in \{0, 1, \dots, \Lambda - \theta - 1\}$), the tagged customer knows that she is in a system with $q = \{\theta, \theta + 1, \dots, \theta + j\}$ customers, with respective probabilities $\mathbb{P}(N = 0), \mathbb{P}(N = 1), \dots, \mathbb{P}(N = j)$, 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:

$$\begin{aligned} \mathbb{E}[Q|\zeta = l] &= \frac{1}{\Lambda - \theta} (\theta + (\theta + (p_H + (1 - p_H)\alpha_l)) + (\theta + 2(p_H + (1 - p_H)\alpha_l)) \\ &\quad + \dots + (\theta + (\Lambda - \theta - 1)(p_H + (1 - p_H)\alpha_l))) \\ &= \frac{1}{\Lambda - \theta} (\theta(\Lambda - \theta) + \frac{(p_H + (1 - p_H)\alpha_l)(\Lambda - \theta - 1)(\Lambda - \theta)}{2}) \\ &= \frac{2\theta + (p_H + (1 - p_H)\alpha_l)(\Lambda - \theta - 1)}{2}. \end{aligned}$$

Then, the tagged customer's expected utility is $\mathbb{E}[U(\alpha'_l, \alpha_l)] = (1 - \alpha'_l)v_L + \alpha'_l(r_L - c(\mathbb{E}[Q|\zeta = l] + 1))$. 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 $r_L - v_L - c(\mathbb{E}[Q|\zeta = l] + 1)$. Recall that $r'_L = r_L - v_L$, and let the root of $r'_L - c(\mathbb{E}[Q|\zeta = l] + 1) = 0$ be

$$\bar{\alpha}_l = \frac{2r'_L - c(2(\theta + 1) + p_H(\Lambda - \theta - 1))}{c(1 - p_H)(\Lambda - \theta - 1)}.$$

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 *long* signal is an equilibrium strategy if, and only if, $0 \in BR(0)$, i.e., $0 \geq \bar{\alpha}_l$, which reduces to $\theta \geq \frac{2r'_L - 2c - cp_H(\Lambda - 1)}{c(2 - p_H)}$. Recall that $\bar{\theta} = \lceil \frac{2r'_L - 2c - cp_H(\Lambda - 1)}{c(2 - p_H)} \rceil$. That is, we have $\alpha_l^* = 0$ whenever $\theta \in \{\bar{\theta}, \dots, \Lambda - 1\}$.
- The strategy $(\alpha_l^* = 1)$ upon receiving a *long* signal is an equilibrium strategy if, and only if, $1 \in BR(1)$, i.e., $1 \leq \bar{\alpha}_l$, which reduces to $\theta \leq \frac{2r'_L - c(\Lambda + 1)}{c}$. Let $\bar{\bar{\theta}} = \lfloor \frac{2r'_L - c(\Lambda + 1)}{c} \rfloor$, we have that $\alpha_l^* = 1$ whenever $\theta \in \{0, \dots, \bar{\bar{\theta}}\}$.
- 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, $\frac{2r'_L - c(\Lambda + 1)}{c} < \theta < \frac{2r'_L - 2c - cp_H(\Lambda - 1)}{c(2 - p_H)}$, that is, $\bar{\bar{\theta}} < \theta < \bar{\theta}$.

■

A.4 Proposition 3 (Fixed Threshold Signaling Mechanism)

PROOF. Based on Lemma 1, we consider three different cases separately: (1) $\theta \leq \bar{\bar{\theta}}$, (2) $\theta \geq \bar{\theta}$, and (3) $\bar{\bar{\theta}} < \theta < \bar{\theta}$.

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

Case 2: $\theta \geq \bar{\theta}$. From Lemma 1, a threshold $\theta \geq \bar{\theta}$ induces a customer equilibrium where low-need customers that receive a short signal join with probability $\alpha_s^* = 1$, and those that receive a *long* signal join with probability $\alpha_l^* = 0$. Based on this, and recalling that all high-need customers always join, we can compute the expected social welfare as follows:

$$\begin{aligned} \Omega(\theta) &= r_H(\theta p_H + \mathbb{E}[N_H]) + r_L\theta(1 - p_H) - \frac{c}{2}(\mathbb{E}[J(\theta)^2] + \mathbb{E}[J(\theta)]) + v_L(\Lambda - \mathbb{E}[J(\theta)]) \\ &= r_H(\theta p_H + \mathbb{E}[N_H]) + r_L\theta(1 - p_H) \\ &\quad - \frac{c}{2}(\theta^2 + 2\theta\mathbb{E}[N_H] + \mathbb{E}[N_H^2] + \theta + \mathbb{E}[N_H]) + v_L(\Lambda - \theta - \mathbb{E}[N_H]). \end{aligned}$$

Since in this case the first θ customers join, and after the threshold θ only high-need customers join, we have that $J(\theta) = \theta + N_H$ is a random variable that represents the total number of customers given threshold θ , where $N_H \sim \mathcal{B}(\Lambda - \theta, p_H)$ is a binomial random variable that represents the total number of high-need customers that join after the threshold θ . From this, we have that $\mathbb{E}[N_H] = (\Lambda - \theta)p_H$, and $\mathbb{E}[N_H^2] = (\Lambda - \theta)p_H + (\Lambda - \theta)(\Lambda - \theta - 1)p_H^2$. Now, consider the following expression:

$$\Omega(\theta) - \Omega(\theta - 1) = (1 - p_H)(r'_L - c(\Lambda p_H + (1 - p_H)\theta)).$$

Notice that $r'_L - c(\Lambda p_H + (1 - p_H)\theta)$ decreases strictly in θ . Based on this, and since we have by construction that $p_H \geq \frac{r'_L}{c\Lambda}$, 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 $\theta \leq \bar{\theta}$). 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 $\theta \leq \bar{\theta}$.

Case 3: $\bar{\theta} < \theta < \bar{\theta}$. From Lemma 1, a threshold $\bar{\theta} < \theta < \bar{\theta}$ induces a customer equilibrium where low-need customers that receive a short signal join $\alpha_s^* = 1$, and those that receive a long signal join with probability $\alpha_l^* = \frac{2r'_L - c(2(\theta+1) + p_H(\Lambda - \theta - 1))}{c(1 - p_H)(\Lambda - \theta - 1)}$. Based on this, and recalling that all high-need customers always join, we can compute the expected social welfare as follows:

$$\begin{aligned} \Omega(\theta) &= r_H(\theta p_H + \mathbb{E}[N_H]) + r_L(\theta(1 - p_H) + \mathbb{E}[N_L]) \\ &\quad - \frac{c}{2}(\mathbb{E}[J(\theta)^2] + \mathbb{E}[J(\theta)]) + v_L(\Lambda - \mathbb{E}[J(\theta)]) \\ &= r_H(\theta p_H + \mathbb{E}[N_H]) + r_L(\theta(1 - p_H) + \mathbb{E}[N] - \mathbb{E}[N_H]) \\ &\quad - \frac{c}{2}(\theta^2 + 2\theta\mathbb{E}[N] + \mathbb{E}[N^2] + \theta + \mathbb{E}[N]) + v_L(\Lambda - \mathbb{E}[J(\theta)]). \end{aligned}$$

Since in this case the first θ customers join, and after the threshold θ high-need customers always join and low-need join with probability α_l^* , we have that $J(\theta) = \theta + N$ is a random variable that represents the total number of customers that join 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 . Moreover, we have that $N = N_H + N_L$, where N_H (respectively, N_L) are random variables that represent the total number of high-need (respectively, low-need) customers that join the system after threshold θ . In particular, $N_H|N \sim \mathcal{B}(N, \frac{p_H}{p_H + (1 - p_H)\alpha_l^*})$, such that $N_H \sim \mathcal{B}(\Lambda - \theta, p_H)$ from the property of conditional binomials. From this, we have that $\mathbb{E}[N_H] = (\Lambda - \theta)p_H$, $\mathbb{E}[N] = (\Lambda - \theta)(p_H + (1 - p_H)\alpha_l^*)$, 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:

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

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

Finally, based on the above cases, 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:

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

We let $g(\bar{\theta}) \triangleq 2(r'_L - c\bar{\theta}) + (\Lambda - \bar{\theta})p_H(2r'_L - c(2(1 + \bar{\theta}) + p_H(\Lambda - \bar{\theta} - 1)))$. We note that the quantity $r'_L - c\bar{\theta} \geq 0 \iff \bar{\theta} \leq \frac{r'_L}{c}$ since $\bar{\theta} \leq q^*$, and that the quantity $2r'_L - c(2(1 + \bar{\theta}) + p_H(\Lambda - \bar{\theta} - 1)) \leq 0 \iff \bar{\theta} \geq \frac{2r'_L - 2c - cp_H(\Lambda - 1)}{c(2 - p_H)}$, since $\bar{\theta} = \lceil \frac{2r'_L - 2c - cp_H(\Lambda - 1)}{c(2 - p_H)} \rceil$. 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 $\alpha_i^*(\bar{\theta}) = 0$ arise in equilibrium. Otherwise, we have that $\theta^* = \bar{\theta} - 1$ and $\alpha_i^*(\bar{\theta} - 1) \in (0, 1)$. ■

A.5 Proposition 4 (No Commitment)

PROOF. We show that the only thresholds that can be sustained in equilibrium are uninformative thresholds. For this, we study the conditions under which the service provider does not have the incentive to implement a θ that deviates from the communicated θ' given customers' best responses to θ' , ($\alpha_s^*(\theta') = 1, \alpha_i^*(\theta') \in [0, 1]$) as described in Lemma 1. For signals to be credible under a communicated θ' , it is required that the expected social welfare, for $\alpha_s^*(\theta') = 1, \alpha_i^*(\theta') \in [0, 1]$, complies with $\Omega(\theta') \geq \Omega(\theta)$ for all possible thresholds θ . Consider the following expression:

$$\begin{aligned} \Omega(\theta') - \Omega(\theta' - 1) = \\ (1 - p_H - (1 - p_H)\alpha_i^*(\theta'))(r'_L - c\Lambda(p_H + (1 - p_H)\alpha_i^*(\theta')) - c(1 - p_H - (1 - p_H)\alpha_i^*(\theta'))\theta'). \end{aligned} \quad (2)$$

First, from (2), we note that if $\theta' \in \{0, \dots, \bar{\theta}, \Lambda\}$, then we have that $\Omega(\theta') = \Omega(\theta' - 1)$ for any threshold θ' . This implies that $\theta' \in \{0, \dots, \bar{\theta}, \Lambda\}$ provide credible signals. However, since $\alpha_s^*(\theta') = 1, \alpha_i^*(\theta') = 1$ in this case, signals do not change customer behavior in comparison to the No-information case - signals are uninformative. Now, from (2), one can show that if $p_H \geq \frac{r'_L}{c\Lambda}$ (which is the case by construction), for any $\alpha_i \in [0, 1]$, we have that $\Omega(\theta') < \Omega(\theta' - 1)$. This means that none of the thresholds $\theta' \in \{\bar{\theta} + 1, \dots, \Lambda - 1\}$ (for which $\alpha_i \in [0, 1]$) can be sustained in equilibrium since their respective signals are not credible — for any of those θ' , the provider has the incentive to implement $\theta = 0$. Overall, since customers are not influenced by signals, 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 signals represent mutual best responses between the customers and the service provider. ■

B Over-Joining Behavior Under Full-Information

Proposition 2 shows that low-need customers join more than is socially optimal, essentially because they ignore the negative externalities that their joining decisions inflict on future arrivals (Naor 1969, Haviv and Oz 2018). In our model, this is particularly the case because low-need customers overcrowd the system to the detriment of high-need customers. To see why, consider the case of only low-need customers ($p_H = 0$). Proposition 2 demonstrates that the socially optimal and equilibrium thresholds coincide ($q_{soc}^* = q^*$), implying no overjoining. This differs from the standard $M/M/1$ queue, where $q_{soc}^* < q^*$ (Naor 1969, Economou 2021).

To see why, for the case of homogeneous customers, consider the expressions for the expected social welfare, $\Omega(n)$, as a function of a joining threshold n (customers join up to position n), in our setting and in the $M/M/1/n$ model:

$$\Omega(n) = rn - c\mathbb{E}[Q] = rn - \sum_{k=1}^n k = rn - \frac{c}{2}(n+1)n, \quad (\text{Our model})$$

$$\Omega(n) = r\lambda_e(n) - c\mathbb{E}[Q] = r\lambda \frac{1 - \rho^n}{1 - \rho^{n+1}} - c \left(\frac{\rho}{1 - \rho} - \frac{(n+1)\rho^{n+1}}{1 - \rho^{n+1}} \right), \quad (M/M/1/n \text{ model})$$

where $\lambda_e(n)$ is the steady-state throughput, under strategy n . In our model $\Omega(n)$ is the expected sum of all customers utilities. In the $M/M/1$ model, it represents the expected utility per time unit. In our model, we have that:

$$\Omega(n) - \Omega(n-1) = (rn - r(n-1)) - \left(\frac{c}{2}(n+1)n - \frac{c}{2}n(n-1) \right) = r - cn > 0 \iff n < \frac{r}{c},$$

which means that if we reduce the threshold (from Naor $q^* = \lfloor \frac{r}{c} \rfloor$), we lose more reward, $(rn - r(n-1))$, than we save on waiting costs, $(\frac{c}{2}(n+1)n - \frac{c}{2}n(n-1))$. This is not necessarily the case in the $M/M/1/n$ model, where it is known that reducing the threshold from Naor can save more on waiting costs than it loses in reward (Naor 1969, Economou 2021). This because the effective arrival rate $\lambda_e(n)$ in the $M/M/1/n$ model is not affected much by reducing the threshold (from Naor) since customers do not always fill up the system up to Naor. Contrary to this, the queue always fills up to Naor in our model due to the “server commences service after all customers make a choice” feature we implement for experimental amenability.

C How to Compute Implied (Expected) Social Welfare

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 social welfare given the participant strategies in our data. Given the design of our experiment, it is natural to compute this at the round-cohort level. We will see that in the *Commit* and *NoCommit* treatments, the computation of welfare 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.

C.1 *NoInfo* Treatment

In this treatment, for each customer $i \in \{1, \dots, \Lambda\}$, in each round $t \in \{1, \dots, T\}$, and cohort $c \in \{1, \dots, C\}$, we observe their strategy $A_{itc} \equiv \{a_{itc}^H \in \{1, 0\}, a_{itc}^L \in \{1, 0\}\}$, where 1 represents the *Join* action and 0 the *Balk* action. Based on this, we compute $\varphi_{tc}^H = \frac{1}{\Lambda} \sum_i a_{itc}^H$ and $\varphi_{tc}^L = \frac{1}{\Lambda} \sum_i a_{itc}^L$, that is, the implied joining proportions (conditional on type), in a particular round and cohort. Note that φ_{tc}^H and φ_{tc}^L 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.

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, $(r_H - c(q+1))p_H\varphi_{tc}^H + (r_L - c(q+1))(1-p_H)\varphi_{tc}^L + v_H p_H(1-\varphi_{tc}^H) + v_L(1-p_H)(1-\varphi_{tc}^L)$, plus the expected utility from time $w-1$ onward, $V_{tc}(q+1, w-1)(p_H\varphi_{tc}^H + (1-p_H)\varphi_{tc}^L) + V_{tc}(q, w-1)(p_H(1-\varphi_{tc}^H) + (1-p_H)(1-\varphi_{tc}^L))$. It follows that:

$$\begin{aligned} V_{tc}(q, w) &= (r_H - c(q+1))p_H\varphi_{tc}^H + (r_L - c(q+1))(1-p_H)\varphi_{tc}^L + v_H p_H(1-\varphi_{tc}^H) + v_L(1-p_H)(1-\varphi_{tc}^L) \\ &\quad + V_{tc}(q+1, w-1)(p_H\varphi_{tc}^H + (1-p_H)\varphi_{tc}^L) + V_{tc}(q, w-1)(p_H(1-\varphi_{tc}^H) + (1-p_H)(1-\varphi_{tc}^L)), \end{aligned}$$

with boundary conditions $V_{tc}(q, 0) = (r_H - c(q+1))p_H\varphi_{tc}^H + (r_L - c(q+1))(1-p_H)\varphi_{tc}^L + v_H p_H(1-\varphi_{tc}^H) + v_L(1-p_H)(1-\varphi_{tc}^L)$ 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 $V_{tc}(0, \Lambda-1)$.

C.2 *FullInfo* Treatment

In this treatment, in each round $t \in \{1, \dots, T\}$, for each cohort $c \in \{1, \dots, C\}$, for each customer $i \in \{1, \dots, \Lambda\}$, we observe their strategy $A_{itc} \equiv \{a_{itc}^H(q) \in \{0, 1\} \ \forall \ q \in \{0, \dots, \Lambda-1\}, a_{itc}^L(q) \in \{0, 1\} \ \forall \ q \in \{0, \dots, \Lambda-1\}\}$, where 1 represents the *Join* action and 0 the *Balk* action. Based on this, we compute $\varphi_{tc}^H(q) = \frac{1}{\Lambda} \sum_i a_{itc}^H(q)$ $\forall \ q \in \{0, \dots, \Lambda-1\}$, $\varphi_{tc}^L(q) = \frac{1}{\Lambda} \sum_i a_{itc}^L(q)$ $\forall \ q \in \{0, \dots, \Lambda-1\}$. That is, the implied joining proportions (conditional on type) for a given queue state in a particular round and cohort. We note that $\varphi_{tc}^H(q)$ and $\varphi_{tc}^L(q)$ 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.

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, $(r_H - c(q+1))p_H\varphi_{tc}^H(q) + (r_L - c(q+1))(1-p_H)\varphi_{tc}^L(q) + v_H p_H(1-\varphi_{tc}^H(q)) + v_L(1-p_H)(1-\varphi_{tc}^L(q))$, plus the expected utility from time $w-1$ onward,

$V_{tc}(q+1, w-1)(p_H\varphi_{tc}^H(q) + (1-p_H)\varphi_{tc}^L(q)) + V_{tc}(q, w-1)(p_H(1-\varphi_{tc}^H(q)) + (1-p_H)(1-\varphi_{tc}^L(q)))$. It follows that:

$$V_{tc}(q, w) = (r_H - c(q+1))p_H\varphi_{tc}^H(q) + (r_L - c(q+1))(1-p_H)\varphi_{tc}^L(q) + v_H p_H(1-\varphi_{tc}^H(q)) + v_L(1-p_H)(1-\varphi_{tc}^L(q)) \\ + V_{tc}(q+1, w-1)(p_H\varphi_{tc}^H(q) + (1-p_H)\varphi_{tc}^L(q)) + V_{tc}(q, w-1)(p_H(1-\varphi_{tc}^H(q)) + (1-p_H)(1-\varphi_{tc}^L(q))),$$

with boundary conditions $V_{tc}(q, 0) = (r_H - c(q+1))p_H\varphi_{tc}^H(q) + (r_L - c(q+1))(1-p_H)\varphi_{tc}^L(q) + v_H p_H(1-\varphi_{tc}^H(q)) + v_L(1-p_H)(1-\varphi_{tc}^L(q))$ 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 $V_{tc}(0, \Lambda - 1)$.

C.3 Commit and NoCommit Treatments

In these treatments, in each round $t \in \{1, \dots, T\}$, for each cohort $c \in \{1, \dots, C\}$, we observe the provider j 's implemented threshold decision θ_{jtc} (*Commit*, *NoCommit*), and the communicated threshold decision θ'_{jtc} (*NoCommit*). For each customer $i \in \{1, \dots, \Lambda\}$, we observe their strategy $A_{itc} \equiv \{a_{itc}^H(\varsigma = \text{short}|\theta_{jtc}) \in \{0, 1\}, a_{itc}^H(\varsigma = \text{long}|\theta_{jtc}) \in \{0, 1\}, a_{itc}^L(\varsigma = \text{short}|\theta_{jtc}) \in \{0, 1\}, a_{itc}^L(\varsigma = \text{long}|\theta_{jtc}) \in \{0, 1\}\}$ (*Commit*), and $A_{itc} \equiv \{a_{itc}^H(\varsigma = \text{short}|\theta'_{jtc}) \in \{0, 1\}, a_{itc}^H(\varsigma = \text{long}|\theta'_{jtc}) \in \{0, 1\}, a_{itc}^L(\varsigma = \text{short}|\theta'_{jtc}) \in \{0, 1\}, a_{itc}^L(\varsigma = \text{long}|\theta'_{jtc}) \in \{0, 1\}\}$ (*NoCommit*), where 1 represents the *Join* action and 0 the *Balk* action. Based on this, in the *Commit* treatment, we compute $\varphi_{tc}^{H, \text{short}} = \frac{1}{\Lambda} \sum_i a_{itc}^H(\varsigma = \text{short}|\theta_{jtc})$, $\varphi_{tc}^{H, \text{long}} = \frac{1}{\Lambda} \sum_i a_{itc}^H(\varsigma = \text{long}|\theta_{jtc})$, $\varphi_{tc}^{L, \text{short}} = \frac{1}{\Lambda} \sum_i a_{itc}^L(\varsigma = \text{short}|\theta_{jtc})$, and $\varphi_{tc}^{L, \text{long}} = \frac{1}{\Lambda} \sum_i a_{itc}^L(\varsigma = \text{long}|\theta_{jtc})$. That is, the implied joining proportions (conditional on type) for a given message in a particular round and cohort. Similarly, in the *NoCommit* treatment, we compute $\varphi_{tc}^{H, \text{short}} = \frac{1}{\Lambda} \sum_i a_{itc}^H(\varsigma = \text{short}|\theta'_{jtc})$, $\varphi_{tc}^{H, \text{long}} = \frac{1}{\Lambda} \sum_i a_{itc}^H(\varsigma = \text{long}|\theta'_{jtc})$, $\varphi_{tc}^{L, \text{short}} = \frac{1}{\Lambda} \sum_i a_{itc}^L(\varsigma = \text{short}|\theta'_{jtc})$, and $\varphi_{tc}^{L, \text{long}} = \frac{1}{\Lambda} \sum_i a_{itc}^L(\varsigma = \text{long}|\theta'_{jtc})$. We note that $\varphi_{tc}^{H, \text{short}}$, $\varphi_{tc}^{H, \text{long}}$, $\varphi_{tc}^{L, \text{short}}$, and $\varphi_{tc}^{L, \text{long}}$ 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.

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, $(r_H - c(q+1))p_H(\sigma_{tc}(q, w)\varphi_{tc}^{H, \text{short}} + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{H, \text{long}}) + (r_L - c(q+1))(1-p_H)(\sigma_{tc}(q, w)\varphi_{tc}^{L, \text{short}} + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{L, \text{long}}) + v_H(p_H(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{H, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{H, \text{long}})) + v_L(1-p_H)(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{L, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{L, \text{long}}))$, plus the expected utility from time $w-1$ onward, $V_{tc}(q+1, w-1)(p_H(\sigma_{tc}(q, w)\varphi_{tc}^{H, \text{short}} + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{H, \text{long}}) + (1-p_H)(\sigma_{tc}(q, w)\varphi_{tc}^{L, \text{short}} + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{L, \text{long}})) + V_{tc}(q, w-1)(p_H(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{H, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{H, \text{long}})) + (1-p_H)(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{L, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{L, \text{long}})))$. It follows that:

$$V_{tc}(q, w) = (r_H - c(q+1))p_H(\sigma_{tc}(q, w)\varphi_{tc}^{H, \text{short}} + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{H, \text{long}}) + (r_L - c(q+1))(1-p_H)(\sigma_{tc}(q, w)\varphi_{tc}^{L, \text{short}} \\ + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{L, \text{long}}) + v_H(p_H(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{H, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{H, \text{long}})) \\ + v_L(1-p_H)(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{L, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{L, \text{long}})) + V_{tc}(q+1, w-1)(p_H(\sigma_{tc}(q, w)\varphi_{tc}^{H, \text{short}} \\ + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{H, \text{long}}) + (1-p_H)(\sigma_{tc}(q, w)\varphi_{tc}^{L, \text{short}} + (1 - \sigma_{tc}(q, w))\varphi_{tc}^{L, \text{long}})) \\ + V_{tc}(q, w-1)(p_H(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{H, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{H, \text{long}})) \\ + (1-p_H)(\sigma_{tc}(q, w)(1 - \varphi_{tc}^{L, \text{short}}) + (1 - \sigma_{tc}(q, w))(1 - \varphi_{tc}^{L, \text{long}}))),$$

with boundary conditions $V_{tc}(q, 0) = (r_H - c(q + 1))p_H(\sigma_{tc}(q, 0)\varphi_{tc}^{H, \text{short}} + (1 - \sigma_{tc}(q, 0))\varphi_{tc}^{H, \text{long}}) + (r_L - c(q + 1))(1 - p_H)(\sigma_{tc}(q, 0)\varphi_{tc}^{L, \text{short}} + (1 - \sigma_{tc}(q, 0))\varphi_{tc}^{L, \text{long}}) + v_H(p_H(\sigma_{tc}(q, 0)(1 - \varphi_{tc}^{H, \text{short}}) + (1 - \sigma_{tc}(q, 0))(1 - \varphi_{tc}^{H, \text{long}})) + v_L(1 - p_H)(\sigma_{tc}(q, 0)(1 - \varphi_{tc}^{L, \text{short}}) + (1 - \sigma_{tc}(q, 0))(1 - \varphi_{tc}^{L, \text{long}}))$ 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 signaling 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.

D Experimental Instructions

D.1 User Interface Snapshots

Here we present snapshots for providers (see Figures 8 - 11) and customers (see Figures 12 - 19) choice elicitation in all treatments. We also present snapshots regarding simulated queue dynamics in the experiment: After all participants have selected their strategies, the computer simulates customer arrivals and any signal communication (if applicable). It then executes each participant's decisions based on their chosen strategies and random realizations (i.e., customer indices and types). In Studies 1 and 2, the simulation is presented as a graphical animation, where each customer arrival and corresponding decision is shown sequentially (see Figures 20 and 21). In Study 3, the simulation is presented in a tabular animation format, where each row is revealed one at a time from top to bottom (see Figures 22, 23, and 24).

Figure 8 Provider Choice Elicitation in *Commit* (Study 1)

Your strategy	☐	☐	☐	☐	☐
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message					

Submit Strategy

(a) Before choice is made

Your strategy	☐	☐	☒	☐	☐
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait	

Submit Strategy

(b) After choice is made

Figure 9 Provider Choice Elicitation in *NoCommit* (Study 1)

Your implemented strategy:	☐	☐	☐	☐	☐
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message					

Your communicated strategy:	☐	☐	☐	☐	☐
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message					

Submit Strategy

(a) Before choice is made

Your implemented strategy:	☐	☐	☐	☐	☒
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message	Long Wait	Long Wait	Long Wait	Long Wait	

Your communicated strategy:	☐	☐	☒	☐	☐
If the number of customers in the waitlist is	3	2	1	0	
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait	


Submit Strategy

(b) After choice is made

Figure 10 Provider Choice Elicitation in *Commit* (Study 3)

(a) Before choice is made

(b) After choice is made

Figure 11 Provider Choice Elicitation in *NoCommit* (Study 3)

(a) Before choice is made

(b) After choice is made

Figure 12 Customer Choice Elicitation in *NoInfo* (Study 1)

If you are H type (70% probability), you will	☐ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☐ Not Join

Submit Strategy

If you are H type (70% probability), you will	☒ Join ☐ Not Join
If you are L type (30% probability), you will	☒ Join ☐ Not Join

Submit Strategy

(a) Before choice is made

(b) After choice is made

Figure 13 Customer Choice Elicitation in *Commit* (Study 1)

Manager's selected strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait

Please select your strategy

	Long Wait	Short Wait
If you are H type (70% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join

[Submit Strategy](#)

(a) Before choice is made

Manager's selected strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait

Please select your strategy

	Long Wait	Short Wait
If you are H type (70% probability), you will	☒ Join ☐ Not Join	☒ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☒ Not Join	☒ Join ☐ Not Join

[Submit Strategy](#)

(b) After choice is made

Figure 14 Customer Choice Elicitation in *NoCommit* (Study 1)

Manager's implemented strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Not Observable			

Manager's communicated strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait

Please select your strategy

	Long Wait	Short Wait
If you are H type (70% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join

[Submit Strategy](#)

(a) Before choice is made

Manager's implemented strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Not Observable			

Manager's communicated strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Short Wait	Short Wait

Please select your strategy

	Long Wait	Short Wait
If you are H type (70% probability), you will	☒ Join ☐ Not Join	☒ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☒ Not Join	☒ Join ☐ Not Join

[Submit Strategy](#)

(b) After choice is made

Figure 15 Customer Choice Elicitation in *CommitBOT* (Study 2)

Manager's (automated bot) selected strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Long Wait	Short Wait

Please select your strategy

	Long Wait	Short Wait
If you are H type (70% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join

[Submit Strategy](#)

(a) Before choice is made

Manager's (automated bot) selected strategy:

If the number of customers in the waitlist is	3	2	1	0
The arriving customer receives message	Long Wait	Long Wait	Long Wait	Short Wait

Please select your strategy

	Long Wait	Short Wait
If you are H type (70% probability), you will	☒ Join ☐ Not Join	☒ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☒ Not Join	☒ Join ☐ Not Join

[Submit Strategy](#)

(b) After choice is made

Figure 16 Customer Choice Elicitation in *FullInfo* Practice Rounds (Study 1,2)

	Waitlist with 3 customer(s)	Waitlist with 2 customer(s)	Waitlist with 1 customer(s)	Waitlist with 0 customer(s)
If you are H type (70% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join	☐ Join ☐ Not Join	☐ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☐ Not Join	☐ Join ☐ Not Join	☐ Join ☐ Not Join	☐ Join ☐ Not Join

[Submit Strategy](#)

(a) Before choice is made

	Waitlist with 3 customer(s)	Waitlist with 2 customer(s)	Waitlist with 1 customer(s)	Waitlist with 0 customer(s)
If you are H type (70% probability), you will	☒ Join ☐ Not Join	☒ Join ☐ Not Join	☒ Join ☐ Not Join	☒ Join ☐ Not Join
If you are L type (30% probability), you will	☐ Join ☒ Not Join	☐ Join ☒ Not Join	☐ Join ☒ Not Join	☐ Join ☒ Not Join

[Submit Strategy](#)

(b) After choice is made

Figure 17 Customer Choice Elicitation in *NoSignal* (Study 3)

If you are "Low-need" (35% probability), you will:

☐ Join
☐ Not Join

If you are "High-need" (65% probability), you will automatically join the waiting list.

[Submit Strategy](#)

(a) Before choice is made

If you are "Low-need" (35% probability), you will:

☒ Join
☐ Not Join

If you are "High-need" (65% probability), you will automatically join the waiting list.

[Submit Strategy](#)

(b) After choice is made

Figure 18 Customer Choice Elicitation in *Commit* (Study 3)

Manager's information strategy

Customer arrives at waiting list of	Customer receives message
0	Short Wait
1	Short Wait
2	Short Wait
3	Short Wait
4	Long Wait
5	Long Wait
6	Long Wait
7	Long Wait

Select your strategy

If you are "Low-need" (35% probability),
If you receive message: You will:

Short Wait ☐ Join
☐ Not Join

Long Wait ☒ Join
☒ Not Join

If you are "High-need" (65% probability), you will automatically join the waiting list.

[Submit Strategy](#)

(a) Before choice is made

Manager's information strategy

Customer arrives at waiting list of	Customer receives message
0	Short Wait
1	Short Wait
2	Short Wait
3	Short Wait
4	Long Wait
5	Long Wait
6	Long Wait
7	Long Wait

Select your strategy

If you are "Low-need" (35% probability),
If you receive message: You will:

Short Wait ☐ Join
☐ Not Join

Long Wait ☒ Join
☒ Not Join

If you are "High-need" (65% probability), you will automatically join the waiting list.

[Submit Strategy](#)

(b) After choice is made

Figure 19 Customer Choice Elicitation in *NoCommit* (Study 3)

Manager's implemented information strategy

Customer arrives at waiting list of	Customer receives message
0	
1	
2	
3	
4	Not Observable
5	
6	
7	

Manager's communicated information strategy

Customer arrives at waiting list of	Customer receives message
0	Short Wait
1	Short Wait
2	Short Wait
3	Short Wait
4	Long Wait
5	Long Wait
6	Long Wait
7	Long Wait

Select your strategy

If you are "Low-need" (35% probability),
If you receive message: You will:

Short Wait ☐ Join
☐ Not Join

Long Wait ☒ Join
☒ Not Join

If you are "High-need" (65% probability), you will automatically join the waiting list.

[Submit Strategy](#)

(a) Before choice is made

Manager's implemented information strategy

Customer arrives at waiting list of	Customer receives message
0	
1	
2	
3	
4	Not Observable
5	
6	
7	

Manager's communicated information strategy

Customer arrives at waiting list of	Customer receives message
0	Short Wait
1	Short Wait
2	Short Wait
3	Short Wait
4	Long Wait
5	Long Wait
6	Long Wait
7	Long Wait

Select your strategy

If you are "Low-need" (35% probability),
If you receive message: You will:

Short Wait ☒ Join
☐ Not Join

Long Wait ☒ Join
☒ Not Join

If you are "High-need" (65% probability), you will automatically join the waiting list.

[Submit Strategy](#)

(b) After choice is made

Figure 20 Queue Simulation in *NoInfo* (Study 1)




Figure 21 Queue Simulation in *Commit*, *NoCommit* (Study 1), and *CommitBOT* (Study 2)

Figure 23 Queue Simulation in *Commit* (Study 3)

[illegible]

Figure 24 Queue Simulation in *NoCommit* (Study 3)

Manager's implemented information strategy		Manager's communicated information strategy		Results (messages are based on the implemented information strategy):									
Customer arrives at waiting list of	Customer receives message	Customer arrives at waiting list of	Customer receives message	Customer	Order of Arrival	Type	Arriving at waiting list of	Received Message	Decision	Position in Waiting List (P)	Service Value (\$)	Waiting Cost (W) = \$40*(P)	Customer Payoff = (\$)(W)
0	Long Wait	0	Short Wait	CC	1	High-need	0	Long Wait	Join	1	\$185	\$40	\$145
1	Long Wait	1	Short Wait	CF	2	High-need	1	Long Wait	Join	2	\$185	\$80	\$105
2	Long Wait	2	Short Wait	CA	3	Low-need	2	Long Wait	Not Join	None	\$0	\$0	\$0
3	Long Wait	3	Short Wait	CH	4	High-need	2	Long Wait	Join	3	\$185	\$120	\$65
4	Long Wait	4	Long Wait	CB	5	Low-need	3	Long Wait	Not Join	None	\$0	\$0	\$0
5	Long Wait	5	Long Wait	CE	6	Low-need	3	Long Wait	Not Join	None	\$0	\$0	\$0
6	Long Wait	6	Long Wait	CG	7	Low-need	3	Long Wait	Not Join	None	\$0	\$0	\$0
7	Long Wait	7	Long Wait	CD	8	High-need	3	Long Wait	Join	4	\$185	\$160	\$25
Manager's Payoff = Average Customer Payoff \$42.50													

D.2 Printed Instructions

In the experiment, participants follow detailed, step-by-step on-screen instructions to ensure full understanding of the task. Additionally, they receive printed handouts that summarize these instructions for easy reference. Below, we present the printed handouts used in the studies. To reduce decision noise and make the conflict of interest more salient in Study 3, the instructions in that study provide participants with a "Decision Support" table that computes payoffs for customers and providers based on the number of customers in queue.

D.2.1 Study 1: *NoInfo* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment that is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, 4 players will take on the role of Customers seeking service, while one player will act as the Manager, responsible for managing the service waitlist. Customers arrive one at a time in a random order, with each customer having an equal chance of arriving 1st, 2nd, 3rd, or 4th. Upon arrival, each customer decides whether to Join or Not Join the waitlist.

In each round, each customer can be either High (H) type (with probability 70%) or Low (L) type (with probability 30%). As a result, a mix of different customer types will typically arrive. Customer types differ in the value they receive from service.

Payoffs

Customers who **join** the waitlist:

- Receive a service with value: $V = \$300$ (if High type), $V = \$200$ (if Low type).
- Incur a waiting cost calculated as $\$50 \times (\text{Position in Waitlist})$: \$50 if they joined 1st, \$100 if they joined 2nd,...
- Receive a total Payoff = $V - \$50 \times (\text{Position in Waitlist})$.

Customers who do **not join** the waitlist:

- Receive service elsewhere, with a total Payoff = \$75.

Manager's Payoff: is calculated as the average customer payoff.

Sequence of Events

Each round of the experiment follows the same sequence:

1. First, Customers define their strategy (i.e., Join or Not Join the waitlist if they happen to be of a particular type) without observing the strategies of other customers.
2. Second, AFTER all Customers selected their strategies, the computer:

Waitlist

Position 4	Position 3	Position 2	Position 1
H: $\$300 - \$200 = \$100$	H: $\$300 - \$150 = \$150$	H: $\$300 - \$100 = \$200$	H: $\$300 - \$50 = \$250$
L: $\$200 - \$200 = \$0$	L: $\$200 - \$150 = \$50$	L: $\$200 - \$100 = \$100$	L: $\$200 - \$50 = \$150$

Manager



Manager Payoff =
Average customer
payoff

Not Join

H,L: $\$75$

- randomly determines the order in which Customers arrive.
- randomly determines the type of each Customer (H with 70% probability, and L with 30% probability).
- simulates the arrivals of Customers, executing their decisions (Join or Not Join) according to their selected strategy.
- calculates the payoff for each Customer and for the Manager.

Structure of the experiment.

The experiment has two parts.

Part A has 10 practice rounds that allow you to familiarize yourself with how the system works. In part A, you will play a version of the game where arriving customers can directly observe the waitlist (such that managers do not need to send information). Throughout part A, you will repeatedly assume the role of customers as well as the role of the manager.

Part B has 40 rounds during which you will play for real money. In part B, customers can **not** directly observe the waitlist. Instead, as described above under “Sequence of Events”, they will define their strategy (i.e., join or not, as an L or H type) without observing the waitlist. For the entire duration of part B of the experiment, you will assume the role of either the Manager or of one of the Customers. At the beginning of each round, 1 Manager and 4 Customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

Your final earnings will equal the average payoff from all 40 rounds of part B of the experiment (each \$1.00 laboratory dollar equals €0.10) plus a €5.00 participation fee.

Good Luck!

D.2.2 Study 1: *Commit* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment that is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, 4 players will take on the role of Customers seeking service, while one player will act as the Manager, responsible for managing the service waitlist. Customers arrive one at a time in a random order, with each customer having an equal chance of arriving 1st, 2nd, 3rd, or 4th. The Manager sends some information about the waitlist. Based on this information, customers decide whether to Join or Not Join upon arrival.

In each round, each customer can be either High (H) type (with probability 70%) or Low (L) type (with probability 30%). As a result, a mix of different customer types will typically arrive. Customer types differ in the value they receive from service.

Payoffs

Customers who **join** the waitlist:

- Receive a service with value: $V = \$300$ (if High type), $V = \$200$ (if Low type).
- Incur a waiting cost calculated as $\$50 \times (\text{Position in Waitlist})$: \$50 if they joined 1st, \$100 if they joined 2nd,...
- Receive a total Payoff = $V - \$50 \times (\text{Position in Waitlist})$.

Customers who do **not join** the waitlist:

- Receive service elsewhere, with a total Payoff = \$75.

Manager's Payoff: is calculated as the average customer payoff.

Sequence of Events

Each round of the experiment follows the same sequence:

1. First, the Manager selects an information strategy X that defines which message an arriving customer receives: 'Short Wait' if the number of customers in the waitlist is less than X , or 'Long Wait' if the number of customers in the waiting list is X or more.
2. Second, AFTER the Manager selects, Customers observe the information strategy X . Then, Customers define their strategy (i.e., Join or Not Join the waitlist if they happen to be of a particular type and receive a particular message) without observing the strategies of other customers.
3. Finally, AFTER all Managers and Customers selected their strategies, the computer:
 - (a) randomly determines the order in which Customers arrive.
 - (b) randomly determines the type of each customer (H with 70% probability, and L with 30% probability).

Waitlist

Position 4	Position 3	Position 2	Position 1
H: $\$300 - \$200 = \$100$	H: $\$300 - \$150 = \$150$	H: $\$300 - \$100 = \$200$	H: $\$300 - \$50 = \$250$
L: $\$200 - \$200 = \$0$	L: $\$200 - \$150 = \$50$	L: $\$200 - \$100 = \$100$	L: $\$200 - \$50 = \$150$

Manager



Manager Payoff =
Average customer
payoff

Not Join

H,L: $\$75$

- (c) simulates the arrival of one Customer at a time, sending a message (Short Wait or Long Wait) according to the Manager's information strategy, and executing Customer decisions (Join or Not Join) according to their selected strategy.
- (d) calculates the payoff for each Customer and for the Manager.

Structure of the experiment.

The experiment has two parts.

Part A has 10 practice rounds that allow you to familiarize yourself with how the system works. In part A, you will play a version of the game where arriving customers can directly observe the waitlist (such that managers do not need to send information). Throughout part A, you will repeatedly assume the role of customers as well as the role of the manager.

Part B has 40 rounds during which you will play for real money. In part B, customers can **not** directly observe the waitlist. Instead, as described above under "Sequence of Events", they will define their strategy (i.e., Join or Not Join, as an L or H type) based on signals ('Short Wait' or 'Long Wait') that follow the manager's *implemented* information policy. For the entire duration of part B of the experiment, you will assume the role of either the Manager or of one of the Customers. At the beginning of each round, 1 Manager and 4 Customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

Your final earnings will equal the average payoff from all 40 rounds of part B of the experiment (each \$1.00 laboratory dollar equals €0.10) plus a €5.00 participation fee.

Good Luck!

D.2.3 Study 1: *NoCommit* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment that is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, 4 players will take on the role of Customers seeking service, while one player will act as the Manager, responsible for managing the service waitlist. Customers arrive one at a time in a random order, with each customer having an equal chance of arriving 1st, 2nd, 3rd, or 4th. The Manager sends some information about the waitlist. Based on this information, customers decide whether to Join or Not Join upon arrival.

In each round, each customer can be either High (H) type (with probability 70%) or Low (L) type (with probability 30%). As a result, a mix of different customer types will typically arrive. Customer types differ in the value they receive from service.

Payoffs

Customers who **join** the waitlist:

- Receive a service with value: $V = \$300$ (if High type), $V = \$200$ (if Low type).
- Incur a waiting cost calculated as $\$50 \times (\text{Position in Waitlist})$: \$50 if they joined 1st, \$100 if they joined 2nd,...
- Receive a total Payoff = $V - \$50 \times (\text{Position in Waitlist})$.

Customers who do **not join** the waitlist:

- Receive service elsewhere, with a total Payoff = \$75.

Manager's Payoff: is calculated as the average customer payoff.

Sequence of Events

Each round of the experiment follows the same sequence:

1. First, the Manager selects an information strategy X that defines which message an arriving customer receives: 'Short Wait' if the number of customers in the waitlist is less than X, or 'Long Wait' if the number of customers in the waiting list is X or more.
2. Second, the Manager selects what information strategy Y to communicate to customers: 'Short Wait' if the number of customers in the waiting list is less than Y, or 'Long Wait' if the number of customers in the waiting list is Y or more. Note: the Manager's *communicated* information strategy Y does not have to be the same as the *implemented* information strategy X that later determines 'Short Wait' or 'Long Wait' messages.
3. Third, AFTER the Manager decides, customers only observe the *communicated* information strategy Y. Then, Customers define their strategy (i.e., Join or Not Join the waitlist if they happen to be of a particular type and receive a particular message) without observing the strategies of other customers.

Waitlist

Position 4	Position 3	Position 2	Position 1
H: $\$300 - \$200 = \$100$	H: $\$300 - \$150 = \$150$	H: $\$300 - \$100 = \$200$	H: $\$300 - \$50 = \$250$
L: $\$200 - \$200 = \$0$	L: $\$200 - \$150 = \$50$	L: $\$200 - \$100 = \$100$	L: $\$200 - \$50 = \$150$

Manager



Manager Payoff =
Average customer
payoff

Not Join

H,L: $\$75$

4. Finally, AFTER all Managers and Customers selected their strategies, the computer:

- randomly determines the order in which Customers arrive.
- randomly determines the type of each customer (H with 70% probability, and L with 30% probability).
- simulates the arrival of one Customer at a time, sending a message (Short Wait or Long Wait) according to the Manager's *implemented* information strategy X, and executing Customer decisions (Join or Not Join) according to their selected strategy.
- calculates the payoff for each Customer and for the Manager.

Structure of the experiment.

The experiment has two parts.

Part A has 10 practice rounds that allow you to familiarize yourself with how the system works. In part A, you will play a version of the game where arriving customers can directly observe the waitlist. Throughout part A, you will repeatedly assume the role of customers as well as the role of the manager.

Part B has 40 rounds during which you will play for real money. In part B, customers can **not** directly observe the waitlist. Instead, as described above under "Sequence of Events", they will define their strategy (i.e., join or not, as an L or H type) based on signals ('Short Wait' or 'Long Wait') that follow the manager's *communicated* information policy. For the entire duration of part B of the experiment, you will assume the role of either the Manager or of one of the Customers. At the beginning of each round, 1 Manager and 4 Customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

Your final earnings will equal the average payoff from all 40 rounds of part B of the experiment (each \$1.00 laboratory dollar equals €0.10) plus a €5.00 participation fee.

Good Luck!

D.2.4 Study 2: *CommitBOT* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment that is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, 4 players will take on the role of Customers seeking service, while an automated bot will act as the Manager, responsible for managing the service waitlist. Customers arrive one at a time in a random order, with each customer having an equal chance of arriving 1st, 2nd, 3rd, or 4th. The Manager (automated bot) sends some information about the waitlist. Based on this information, customers decide whether to Join or Not Join upon arrival. In each round, each customer can be either High (H) type (with probability 70%) or Low (L) type (with probability 30%). As a result, a mix of different customer types will typically arrive. Customer types differ in the value they receive from service.

Payoffs

Customers who **join** the waitlist:

- Receive a service with value: $V = \$300$ (if High type), $V = \$200$ (if Low type).
- Incur a waiting cost calculated as $\$50 \times (\text{Position in Waitlist})$: \$50 if they joined 1st, \$100 if they joined 2nd,...
- Receive a total Payoff = $V - \$50 \times (\text{Position in Waitlist})$.

Customers who do **not join** the waitlist:

- Receive service elsewhere, with a total Payoff = \$75.

Manager's (automated bot's) Payoff: is calculated as the average customer payoff.

Sequence of Events

Each round of the experiment follows the same sequence:

1. First, the Manager (automated bot) selects an information strategy X that defines which message an arriving customer receives: 'Short Wait' if the number of customers in the waitlist is less than X, or 'Long Wait' if the number of customers in the waiting list is X or more.
2. Second, AFTER the Manager (automated bot) selects, Customers observe the information strategy X. Then, Customers define their strategy (i.e., Join or Not Join the waitlist if they happen to be of a particular type and receive a particular message) without observing the strategies of other customers.
3. Finally, AFTER the Manager (automated bot) and all Customers selected their strategies, the computer:
 - (a) randomly determines the order in which Customers arrive.
 - (b) randomly determines the type of each customer (H with 70% probability, and L with 30% probability).
 - (c) simulates the arrival of one Customer at a time, sending a message (Short Wait or Long Wait) according to the Manager's (automated bot's) information strategy, and executing Customer decisions (Join or Not Join) according to their selected strategy.

Waitlist

Position 4	Position 3	Position 2	Position 1
H: $\$300 - \$200 = \$100$	H: $\$300 - \$150 = \$150$	H: $\$300 - \$100 = \$200$	H: $\$300 - \$50 = \$250$
L: $\$200 - \$200 = \$0$	L: $\$200 - \$150 = \$50$	L: $\$200 - \$100 = \$100$	L: $\$200 - \$50 = \$150$

Manager



Manager Payoff =
Average customer
payoff

Not Join

H,L: $\$75$

(d) calculates the payoff for each Customer and for the Manager (automated bot).

Structure of the experiment.

The experiment has two parts.

Part A has 10 practice rounds that allow you to familiarize yourself with how the system works. In part A, you will play a version of the game where arriving customers can directly observe the waitlist (such that the Manager (automated bot) do not need to send information).

Part B has 40 rounds during which you will play for real money. In part B, customers can **not** directly observe the waitlist. Instead, as described above under “Sequence of Events”, they will define their strategy (i.e., Join or Not Join, as an L or H type) based on signals (‘Short Wait’ or ‘Long Wait’) that follow the manager’s (automated bot’s) *implemented* information policy. At the beginning of each round, 4 Customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

Your final earnings will equal the average payoff from all 40 rounds of part B of the experiment (each \$1.00 laboratory dollar equals €0.10) plus a €5.00 participation fee.

Good Luck!

D.2.5 Study 3: *NoInfo* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment on decision making. The experiment is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, one of you will assume the role of the Manager of a firm that offers services to customers, and 8 of you will assume the role of Customers. When arriving to the market, customers need to decide whether or not to join a waiting list to receive service. In each round, each customer can be Low-need (with 35% probability) or High-need (with 65% probability). Low-need customers can decide whether or not to join a waiting list to receive service. High-need customers always (automatically) join the waiting list. The Manager answers questions about the decisions of customers.

Detailed description

1. At the beginning of each round, customers define their strategy (i.e., Join or Not Join the waiting list if they happen to be a Low-need customer) without observing the strategies of other customers.
2. AFTER all customers input their strategies, the computer randomly determines the order in which customers arrive to the market. Then it randomly determines the type of each customer (Low-need with 35% probability, and High-need with 65% probability). Finally it implements decisions based on customers' strategies.
3. Given the order of arrivals and customers' decisions, the computer generates a waiting list.
4. Service Value: customers receive \$185 if they join the waiting list, and \$0 otherwise.
5. Waiting Cost: customers who join the waiting list have a waiting cost equal to \$40 TIMES their position in the waiting list.
6. Customers' Payoff: is calculated as the Service Value MINUS the Waiting Cost.
7. Manager's Payoff: is calculated as the average customer payoff.

Your role and your group

For the entire duration of the experiment, you will assume the role of either the Manager or of one of the Customers. At the beginning of each round, 1 manager and 8 customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

You will play 40 experimental rounds for real money. Your final earnings will equal the average payoff from all your rounds (each \$1.00 laboratory dollar equals €0.50) plus €7.00 participation fee.

Decision Support

Customers have a negative payoff when their position in the waiting list is greater than 4. Also, the average customer payoff (and thus the Manager's payoff) is maximized when only 4 customers join the waiting list.

Customer Position in Waiting List	Customer Payoff	Number of Customers that join Waiting List	Manager's Payoff = Average Customer Payoff
None (Not Join)	\$0	0	\$0.00
1	\$145	1	\$18.13
2	\$105	2	\$31.25
3	\$65	3	\$39.38
4	\$25	4	\$42.50
5	-\$15	5	\$40.63
6	-\$55	6	\$33.75
7	-\$95	7	\$21.88
8	-\$135	8	\$5.00

Good Luck!

D.2.6 Study 3: *Commit* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment on decision making. The experiment is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, one of you will assume the role of the Manager of a firm that offers services to customers, and 8 of you will assume the role of Customers. When arriving to the market, customers need to decide whether or not to join a waiting list to receive service. In each round, each customer can be Low-need (with 35% probability) or High-need (with 65% probability). Low-need customers can decide whether or not to join a waiting list to receive service. High-need customers always (automatically) join the waiting list. The Manager sends some information about the waiting list. Based on this information, customers decide whether to Join or Not Join the waiting list.

Detailed description

1. At the beginning of each round the Manager selects an information strategy X that defines which message an arriving customer receives: 'Short Wait' if the number of customers in the waiting list is less than X, or 'Long Wait' if the number of customers in the waiting list is X or more.
2. AFTER the Manager decides, the information strategy is communicated to all customers. Customers need to define their strategies (i.e., Join or Not Join the waiting list if they happen to be a Low-need customer) for a given message that they could receive.
3. AFTER all customers input their strategies, the computer randomly determines the order in which customers arrive to the market. Then it randomly determines the type of each customer (Low-need with 35% probability, and High-need with 65% probability). Finally it implements all participants' strategies to generate the waiting list.
4. Service Value: customers receive \$185 if they join the waiting list, and \$0 otherwise.

5. Waiting Cost: customers who join the waiting list have a waiting cost equal to \$40 TIMES their position in the waiting list.
6. Customers' Payoff: is calculated as the Service Value MINUS the Waiting Cost.
7. Manager's Payoff: is calculated as the average customer payoff.

Your role and your group

For the entire duration of the experiment, you will assume the role of either the Manager or of one of the Customers. At the beginning of each round, 1 manager and 8 customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

You will play 40 experimental rounds for real money. Your final earnings will equal the average payoff from all your rounds (each \$1.00 laboratory dollar equals €0.50) plus €7.00 participation fee.

Decision Support

Customers have a negative payoff when their position in the waiting list is greater than 4. Also, the average customer payoff (and thus the Manager's payoff) is maximized when only 4 customers join the waiting list.

Customer Position in Waiting List	Customer Payoff	Number of Customers that join Waiting List	Manager's Payoff = Average Customer Payoff
None (Not Join)	\$0	0	\$0.00
1	\$145	1	\$18.13
2	\$105	2	\$31.25
3	\$65	3	\$39.38
4	\$25	4	\$42.50
5	-\$15	5	\$40.63
6	-\$55	6	\$33.75
7	-\$95	7	\$21.88
8	-\$135	8	\$5.00

Good Luck!

D.2.7 Study 3: *NoCommit* Treatment

Instructions

You receive detailed on-screen instructions at the beginning of the session. This handout is simply for your convenience in case you wish to revisit the instructions at any time during the experiment. You are about to participate in an experiment on decision making. The experiment is designed to learn about your decision making, not to test your knowledge. All individual responses are completely confidential and anonymous. If you have any question, feel free to raise your hand.

Task Overview

In each of several rounds, one of you will assume the role of the Manager of a firm that offers services to customers, and 8 of you will assume the role of Customers. When arriving to the market, customers need to decide whether or not to join a waiting list to receive service. In each round, each customer can be Low-need (with 35% probability) or High-need (with 65% probability). Low-need customers can decide whether or not

to join a waiting list to receive service. High-need customers always (automatically) join the waiting list. The Manager sends some information about the waiting list. Based on this information, customers decide whether to Join or Not Join the waiting list.

Detailed description

1. At the beginning of each round the Manager selects an information strategy X that defines which message an arriving customer receives: 'Short Wait' if the number of customers in the waiting list is less than X, or 'Long Wait' if the number of customers in the waiting list is X or more.
2. The *implemented* information strategy X is NOT directly observable by customers but the Manager is able to communicate the information strategy to customers. Specifically, the Manager selects what information strategy Y to communicate to customers: 'Short Wait' if the number of customers in the waiting list is less than Y, or 'Long Wait' if the number of customers in the waiting list is Y or more. Note: the Manager's *communicated* information strategy Y does not have to be the same as the *implemented* information strategy X that later determines 'Short Wait' or 'Long Wait' messages.
3. AFTER the Manager decides, customers only observe the *communicated* information strategy Y. Customers need to define their strategies (i.e., Join or Not Join the waiting list if they happen to be a Low-need customer) for a given message that they could receive. Recall that messages are based on the *implemented* information strategy X.
4. AFTER all customers input their strategies, the computer randomly determines the order in which customers arrive to the market. Then it randomly determines the type of each customer (Low-need with 35% probability, and High-need with 65% probability). Finally, it implements all customers' joining strategies, and the Manager's *implemented* information strategy X, to generate the waiting list. Note: at the end of each round, customers know what was the implemented information strategy X, and if the Manager communicated an information strategy Y different from the implemented X.
5. Service Value: customers receive \$185 if they join the waiting list, and \$0 otherwise.
6. Waiting Cost: customers who join the waiting list have a waiting cost equal to \$40 TIMES their position in the waiting list.
7. Customers' Payoff: is calculated as the Service Value MINUS the Waiting Cost.
8. Manager's Payoff: is calculated as the average customer payoff.

Your role and your group

For the entire duration of the experiment, you will assume the role of either the Manager or of one of the Customers. At the beginning of each round, 1 manager and 8 customers will be randomly matched. You will interact with different participants during the experiment and you will never get to know who you are playing with in a given round.

Your earnings from the experiment

You will play 40 experimental rounds for real money. Your final earnings will equal the average payoff from all your rounds (each \$1.00 laboratory dollar equals €0.50) plus €7.00 participation fee.

Decision Support

Customers have a negative payoff when their position in the waiting list is greater than 4. Also, the average customer payoff (and thus the Manager's payoff) is maximized when only 4 customers join the waiting list.

Customer Position in Waiting List	Customer Payoff
None (Not Join)	\$0
1	\$145
2	\$105
3	\$65
4	\$25
5	-\$15
6	-\$55
7	-\$95
8	-\$135

Number of Customers that join Waiting List	Manager's Payoff = Average Customer Payoff
0	\$0.00
1	\$18.13
2	\$31.25
3	\$39.38
4	\$42.50
5	\$40.63
6	\$33.75
7	\$21.88
8	\$5.00

Good Luck!

D.3 Comprehension Quizzes

Participants start reading instructions about either *FullInfo* (Study 1 and 2), or *NoInfo* (Study 3). They then completed Part A of the Quiz (see below), which assessed their general understanding of system parameters, queue dynamics, and the role of random events. All questions had to be answered correctly in order to proceed. After this, participants moved to some practice rounds of *FullInfo* (Study 1 and 2), or *NoInfo* (Study 3). Participants allocated to the *NoInfo* treatment (in Study 3) proceeded to the actual experiment. Participants allocated to the *NoInfo* (Study 1), *Commit* (Studies 1 and 3), *NoCommit* (Studies 1 and 3), or *CommitBOT* (Study 2) treatments received additional instructions tailored to their specific condition. They subsequently completed Part B of the Quiz, which was also treatment-specific. As before, participants were required to answer all questions correctly to proceed to the main experiment. Quizzes are presented to participants, irrespective of their role. *Note:* In the quizzes below, participants were shown the actual numerical values relevant to their study, rather than the algebraic notation used here (e.g., r_H , or expressions such as $(c + 20)$). Correct answers are presented in bold.

D.3.1 Part A Quiz (Study 1 and 2)

[Q1] H type customers receive a payoff of (select one):

- $r_H - c \times$ **(Position Waitlist) when they join, and v_H when they do not join**
- $r_H - (c + 20) \times$ (Position Waitlist) when they join, and v_H when they do not join
- $r_H - c \times$ (Position Waitlist) when they join, and $v_L + 30$ when they do not join

[Q2] L type customers receive a payoff of (select one):

- $r_L - (c + 20) \times$ (Position Waitlist) when they join, and v_H when they do not join
- $r_H - c \times$ (Position Waitlist) when they join, and v_H when they do not join
- $r_L - c \times$ **(Position Waitlist) when they join, and v_L when they do not join**

[Q3] When a customer selects a strategy (Join or Not Join), they:

- ...know their order of arrival and the strategies of other customers
- ...know their order of arrival, but not the strategies of other customers
- **...do not know their order of arrival nor the strategies of other customers**

[Q4] A customer who arrives 2nd in one round, will arrive 2nd in every round.

- Yes
- **No, the computer will randomly determine the arrival order at the beginning of each round**

[Q5] Select the correct option.

- A customer who is H type in one round, will be a H type customer in every round.
- **The computer will randomly determine the type of each customer at the beginning of each round: L type ($100 - p_H\%$ probability), and H type ($p_H\%$ probability)**
- The computer will randomly determine the type of each customer at the beginning of each round: L type (50% probability), and H type (50% probability)

[Q6] In a given round, either all customers are H type, or all customers are L type.

- Yes
- **No, the computer will randomly determine the type of each customer in each round. Hence, a mix of different customer types will typically arrive**

D.3.2 Part A Quiz (Study 3)

[Q1] If a customer's strategy is to NOT Join the waiting list as a Low-need customer, the customer:

- ...will definitely not join the waiting list
- **...might still join the waiting list if the computer randomly assigns High-need**

[Q2] Select the correct option:

- Customers who join the waiting list get a Payoff equal to $r_L - 50 \times (\text{Position in Waiting List})$. Customers who do not join get a Payoff of v_L
- Customers who join the waiting list get a Payoff equal to $\text{\$service_value} - c \times (\text{Position in Waiting List})$. Customers who do not join get a Payoff of r_L
- **Customers who join the waiting list get a Payoff equal to $r_L - c \times (\text{Position in Waiting List})$. Customers who do not join get a Payoff of v_L**

[Q3] If a customer joins the waiting list, they always earn a positive payoff:

- Yes
- **No, the payoff can be negative if the wait cost is large**

[Q4] When a customer selects a strategy (Join or Not Join waiting list as a Low-need customer), they:

- ...know their order of arrival and the strategies of other customers
- ...know their order of arrival, but not the strategies of other customers
- **...do not know their order of arrival nor the strategies of other customers**

[Q5] A customer who arrives 5th in one round, will arrive 5th in every round:

- Yes
- **No, the computer will randomly determine the arrival order at the beginning of each round**

[Q6] A customer who is a High-need customer in one round, will be a High-need customer in every round:

- Yes
- **No, the computer will randomly determine the type of each customer at the beginning of each round: Low-need ($100 - p_H\%$ probability), and High-need ($p_H\%$ chance)**
- No, the computer will randomly determine the type of each customer at the beginning of each round: Low-need (50% probability), and High-need (50% chance)

[Q7] In a given round, either all customers are High-need, or all customers are Low-need:

- Yes
- **No, the computer will randomly determine the type of each customer in each round. Hence, a mix of different customer types will typically arrive**

[Q8] The more customers join the waiting list, the higher the Manager's payoff (customers' average payoff):

- Yes
- **No, the Manager payoff is at its maximum when only 4 customers join the waiting list**

D.3.3 Part B Quiz: *NoInfo* (Study 1)

[Q1] When a customer selects a strategy, this strategy specifies what decision (Join or Not Join) the computer will implement:

- ...based on the customer type (H or L) and the length of the waitlist upon arrival
- **...based on the customer type (H or L), but not on the length of the waitlist upon arrival (which the customer cannot observe in part B)**
- ...based on the length of the waitlist upon arrival, but not on the customer type (H or L)

D.3.4 Part B Quiz: *Commit* (Study 1)

[Q1] When a customer selects a strategy, this strategy specifies what decision (Join or Not Join) the computer will implement:

- ...based on the customer type (H or L) and the length of the waitlist upon arrival
- **...based on the customer type (H or L) and the received message (Short or Long Wait) upon arrival**
- ...based on the received message (Short or Long Wait) upon arrival, but not on the customer type (H or L)

[Q2] Imagine that the Manager decides to send the 'Long Wait' message if the waitlist has 2 or more customers. Then if you are the 3rd customer to arrive, you will:

- Receive the 'Long Wait' message for sure
- **Receive the 'Long Wait' message only when both previous customers joined the waitlist, but the 'Short Wait' message if at least one of them did not join**

D.3.5 Part B Quiz: *NoCommit* (Study 1)

[Q1] When a customer selects a strategy, this strategy specifies what decision (Join or Not Join) the computer will implement:

- ...based on the customer type (H or L) and the length of the waitlist upon arrival
- **...based on the customer type (H or L) and the received message (Short or Long Wait) upon arrival**
- ...based on the received message (Short or Long Wait) upon arrival, but not on the customer type (H or L)

[Q2] Imagine that the Manager decides to send the 'Long Wait' message if the waitlist has 2 or more customers. Then if you are the 3rd customer to arrive, you will:

- Receive the 'Long Wait' message for sure

- **Receive the ‘Long Wait’ message only when both previous customers joined the wait-list, but the ‘Short Wait’ message if at least one of them did not join**

[Q3] Are the Manager’s “implemented information strategy” and the “communicated information strategy” necessarily the same?

- Yes
- **No, the Manager can choose to communicate a different information strategy**

[Q4] Will customers, at the end of each round, know if the Manager communicated an information strategy different from the implemented strategy?

- No
- **Yes, customers will know this when presented with the results at the end of each round**

D.3.6 Part B Quiz: *CommitBOT* (Study 2)

[Q1] When a customer selects a strategy, this strategy specifies what decision (Join or Not Join) the computer will implement:

- ...based on the customer type (H or L) and the length of the waitlist upon arrival
- **...based on the customer type (H or L) and the received message (Short or Long Wait) upon arrival**
- ...based on the received message (Short or Long Wait) upon arrival, but not on the customer type (H or L)

[Q2] Imagine that the Manager (automated Bot) decides to send the ‘Long Wait’ message if the waitlist has 2 or more customers. Then if you are the 3rd customer to arrive, you will (select one):

- Receive the ‘Long Wait’ message for sure
- **Receive the ‘Long Wait’ message only when both previous customers joined the wait-list, but the ‘Short Wait’ message if at least one of them did not join**

[Q3] Another participant will play the role of the Manager:

- True
- **False, an automated Bot will play the role of the Manager**

D.3.7 Part B Quiz: *Commit* (Study 3)

[Q1] Imagine that the Manager decides to send the ‘Long Wait’ message if the list has 5 or more customers. Then if you are the 7th customer to arrive to the market, will you receive the ‘Long Wait’ message for sure?

- Yes
- **No, other customers that arrive before may not join the waiting list so the list might still be shorter than 5 when you arrive**

[Q2] If a customer’s strategy is to Join the waiting list when receiving the ‘Short Wait’ message, and to NOT Join when receiving the ‘Long Wait’ message, then ...

- ...the customer will join the waiting list only if he/she happens to be of High-need
- **...the customer will join the waiting list if he/she is High-need, or if he/she is Low-need and receives a 'Short Wait' message**

D.3.8 Part B Quiz: *NoCommit* (Study 3)

[Q1] Imagine that the Manager decides to send the 'Long Wait' message if the list has 5 or more customers. Then if you are the 7th customer to arrive to the market, will you receive the 'Long Wait' message for sure?

- Yes
- **No, other customers that arrive before may not join the waiting list so the list might still be shorter than 5 when you arrive**

[Q2] If a customer's strategy is to Join the waiting list when receiving the 'Short Wait' message, and to NOT Join when receiving the 'Long Wait' message, then ...

- ...the customer will join the waiting list only if he/she happens to be of High-need
- **...the customer will join the waiting list if he/she is High-need, or if he/she is Low-need and receives a 'Short Wait' message**

[Q3] Are the Manager's 'implemented information strategy' and the 'communicated information strategy' necessarily the same?

- Yes
- **No, the Manager can choose to communicate a different information strategy**

[Q4] Will customers, at the end of each round, know if the Manager communicated an information strategy different from the implemented one?

- No
- **Yes, customers will know this when presented with the results at the end of each round**

E Experimental Results

E.1 Study 1

E.1.1 Customer Choices We examine aggregate patterns in customers' joining behavior. Table 2 presents comparisons between average customer joining probabilities and relevant theoretical benchmarks. Since high-need customers almost always join regardless of treatment, we focus on low-need customers to assess treatment effects. Table 3 compares average joining probabilities for low-need customers between treatments. Table 4 presents logistic regression results for low-need customers' joining probabilities across treatments. Finally, Figures 25–27 show average (across rounds) joining probabilities for low-need customers at the individual participant level in each treatment.

Table 2 T-Tests: Customer Joining Behavior vs. Theoretical Benchmarks in Study 1

Type	Treatment	Signal	Threshold θ'	Avg Prob	Theory/Benchmark	t(df)	p-value	Bootstrap p-value
High-need	<i>NoInfo</i>	–	–	0.998	1.000	$t(7) = -1.52$	0.085	0.028
High-need	<i>Commit</i>	Short Wait	All	0.993	1.000	$t(7) = -1.42$	0.099	0.027
		Long Wait	All	0.983	1.000	$t(7) = -2.90$	0.011	< 0.001
		Short Wait	1	0.979	1.000	$t(3) = -1.00$	0.195	0.050
		Short Wait	2	0.994	1.000	$t(7) = -1.08$	0.157	0.071
		Short Wait	3	0.994	1.000	$t(7) = -1.17$	0.140	0.069
		Short Wait	4	0.994	1.000	$t(6) = -1.00$	0.178	0.061
		Long Wait	0	0.975	1.000	$t(4) = -1.00$	0.187	0.056
		Long Wait	1	0.982	1.000	$t(3) = -1.66$	0.097	0.063
		Long Wait	2	0.976	1.000	$t(7) = -2.96$	0.010	0.004
		Long Wait	3	0.985	1.000	$t(7) = -2.08$	0.038	0.019
High-need	<i>NoCommit</i>	Short Wait	All	1.000	1.000	–	–	–
		Long Wait	All	0.992	1.000	$t(6) = -1.46$	0.097	0.028
		Short Wait	1–4	1.000	1.000	–	–	–
		Long Wait	0	0.916	1.000	$t(2) = -1.00$	0.211	0.038
		Long Wait	1	1.000	1.000	–	–	–
		Long Wait	2	0.993	1.000	$t(6) = -1.51$	0.090	0.026
		Long Wait	3	0.970	1.000	$t(6) = -1.50$	0.091	0.030
Low-need	<i>NoInfo</i>	–	–	0.512	1.000	$t(7) = -10.98$	< 0.001	< 0.001
		–	–	0.512	0.500	$t(7) = 0.27$	0.397	0.399
		–	–	0.512	0.000	$t(7) = 11.52$	< 0.001	< 0.001
Low-need	<i>Commit</i>	Short Wait	All	0.914	1.000	$t(7) = -2.29$	0.027	0.006
		Long Wait	All	0.100	0.000	$t(7) = 3.23$	0.007	0.004
		Short Wait	1	0.937	1.000	$t(3) = -1.00$	0.196	0.054
		Short Wait	2	0.967	1.000	$t(7) = -1.96$	0.045	0.014
		Short Wait	3	0.927	1.000	$t(7) = -2.49$	0.021	0.004
		Short Wait	4	0.685	1.000	$t(6) = -6.25$	< 0.001	0.004
		Long Wait	0	0.528	1.000	$t(4) = -6.97$	0.001	0.010
		Long Wait	1	0.220	0.000	$t(3) = 5.00$	0.008	0.062
		Long Wait	2	0.067	0.000	$t(7) = 2.87$	0.012	0.006
		Long Wait	3	0.056	0.000	$t(7) = 3.35$	0.006	0.005
Low-need	<i>NoCommit</i>	Short Wait	All	0.891	1.000	$t(6) = -2.08$	0.041	0.013
		Long Wait	All	0.049	1.000	$t(6) = -43.52$	< 0.001	< 0.001
		Short Wait	1	0.975	1.000	$t(4) = -1.00$	0.187	0.060
		Short Wait	2	0.966	1.000	$t(6) = -2.86$	0.014	0.008
		Short Wait	3	0.794	1.000	$t(6) = -3.02$	0.011	0.012
		Short Wait	4	0.536	1.000	$t(4) = -4.47$	0.005	0.014
		Long Wait	0	0.229	1.000	$t(2) = -6.64$	0.011	0.038
		Long Wait	1	0.191	1.000	$t(4) = -22.25$	< 0.001	< 0.001
		Long Wait	2	0.044	1.000	$t(6) = -54.88$	< 0.001	< 0.001
		Long Wait	3	0.102	1.000	$t(6) = -15.64$	< 0.001	0.019

Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted joining probabilities for each customer type, or based on benchmarks of interest. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates. Constant values resulted in omitted t-stats.

Table 3 T-Tests: Treatment Effects, Low-Need Customer Joining Behavior in Study 1

Comparison	Signal	Threshold	Avg Prob	t(df)	p-value	Bootstrap p-value	Welch p-value	Welch Bootstrap p-value
<i>Commit vs NoInfo</i>	Short Wait	All	0.914 vs 0.512	t(14) = 6.92	< 0.001	< 0.001	< 0.001	< 0.001
<i>Commit vs NoInfo</i>	Long Wait	All	0.100 vs 0.512	t(14) = -7.60	< 0.001	< 0.001	< 0.001	< 0.001
<i>Commit vs NoInfo</i>	Short Wait	1	0.937 vs 0.512	t(10) = 5.54	< 0.001	< 0.001	< 0.001	0.002
<i>Commit vs NoInfo</i>	Short Wait	2	0.967 vs 0.512	t(14) = 9.61	< 0.001	< 0.001	< 0.001	< 0.001
<i>Commit vs NoInfo</i>	Short Wait	3	0.927 vs 0.512	t(14) = 7.83	< 0.001	< 0.001	< 0.001	< 0.001
<i>Commit vs NoInfo</i>	Short Wait	4	0.685 vs 0.512	t(13) = 2.60	0.011	0.011	0.011	0.012
<i>Commit vs NoInfo</i>	Long Wait	0	0.528 vs 0.512	t(11) = 0.20	0.418	0.418	0.423	0.408
<i>Commit vs NoInfo</i>	Long Wait	1	0.220 vs 0.512	t(10) = -4.65	0.001	0.002	< 0.001	0.002
<i>Commit vs NoInfo</i>	Long Wait	2	0.067 vs 0.512	t(14) = -8.85	< 0.001	< 0.001	< 0.001	< 0.001
<i>Commit vs NoInfo</i>	Long Wait	3	0.056 vs 0.512	t(14) = -9.58	< 0.001	< 0.001	0.011	0.012
<i>NoCommit vs NoInfo</i>	Short Wait	All	0.891 vs 0.512	t(13) = 5.56	< 0.001	< 0.001	< 0.001	< 0.001
<i>NoCommit vs NoInfo</i>	Long Wait	All	0.049 vs 0.512	t(13) = -8.92	< 0.001	< 0.001	< 0.001	0.001
<i>NoCommit vs NoInfo</i>	Short Wait	1	0.975 vs 0.512	t(11) = 9.08	< 0.001	< 0.001	< 0.001	< 0.001
<i>NoCommit vs NoInfo</i>	Short Wait	2	0.966 vs 0.512	t(13) = 9.89	< 0.001	< 0.001	< 0.001	< 0.001
<i>NoCommit vs NoInfo</i>	Short Wait	3	0.794 vs 0.512	t(13) = 3.57	0.002	0.002	0.003	0.003
<i>NoCommit vs NoInfo</i>	Short Wait	4	0.536 vs 0.512	t(11) = 0.25	0.403	0.403	0.417	0.414
<i>NoCommit vs NoInfo</i>	Long Wait	0	0.229 vs 0.512	t(9) = -2.27	0.009	0.011	0.060	0.037
<i>NoCommit vs NoInfo</i>	Long Wait	1	0.191 vs 0.512	t(11) = -5.58	< 0.001	< 0.001	< 0.001	< 0.001
<i>NoCommit vs NoInfo</i>	Long Wait	2	0.044 vs 0.512	t(13) = -9.27	< 0.001	< 0.001	< 0.001	< 0.001
<i>NoCommit vs NoInfo</i>	Long Wait	3	0.102 vs 0.512	t(13) = -5.72	< 0.001	< 0.001	< 0.001	0.002
<i>Commit vs NoCommit</i>	Short Wait	All	0.914 vs 0.891	t(13) = 0.36	0.360	0.369	0.363	0.351
<i>Commit vs NoCommit</i>	Long Wait	All	0.100 vs 0.049	t(13) = 1.29	0.109	0.111	0.103	0.109
<i>Commit vs NoCommit</i>	Short Wait	1	0.937 vs 0.975	t(7) = -0.61	0.281	NA	0.304	0.203
<i>Commit vs NoCommit</i>	Short Wait	2	0.967 vs 0.966	t(13) = 0.05	0.478	0.478	0.477	0.490
<i>Commit vs NoCommit</i>	Short Wait	3	0.927 vs 0.794	t(13) = 1.89	0.040	0.040	0.054	0.052
<i>Commit vs NoCommit</i>	Short Wait	4	0.685 vs 0.536	t(10) = 1.42	0.092	0.092	0.122	0.143
<i>Commit vs NoCommit</i>	Long Wait	0	0.528 vs 0.229	t(6) = 2.42	0.026	0.025	0.051	0.035
<i>Commit vs NoCommit</i>	Long Wait	1	0.220 vs 0.191	t(7) = 0.51	0.311	NA	0.314	0.305
<i>Commit vs NoCommit</i>	Long Wait	2	0.067 vs 0.044	t(13) = 0.75	0.232	0.231	0.227	0.225
<i>Commit vs NoCommit</i>	Long Wait	3	0.056 vs 0.102	t(13) = -0.82	0.214	0.234	0.233	0.186

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 4 Logistic Regressions: Customer Strategies in *NoInfo*, *Commit*, and *NoCommit* in Study 1

	<i>NoInfo</i> P(Join) (1a)	<i>Commit & NoCommit</i> P(Join Short Wait) (1b) (2b) (3b)			<i>Commit & NoCommit</i> P(Join Long Wait) (1c) (2c) (3c)		
(Intercept)	0.99** (0.40)	9.08*** (1.05)	8.90** (0.81)	8.79*** (0.91)	-1.11 (0.99)	-0.65 (0.60)	-0.81 (0.62)
<i>NoCommit</i>	- -	- -	- -	0.39 (1.41)	- -	- -	-0.37 (0.37)
<i>Communicated Threshold</i>	- -	-1.55*** (0.25)	-1.75*** (0.26)	-1.54*** (0.24)	-1.02*** (0.25)	-1.16*** (0.15)	-1.02*** (0.25)
<i>NoCommit * Communicated Threshold</i>	- -	- -	- -	-0.22 (0.36)	- -	- -	-0.14 (0.28)
Round	0.00 (0.00)	-0.01 (0.01)	0.00 (0.01)	0.01 (0.01)	-0.01 (0.01)	-0.01 (0.02)	-0.01 (0.01)
Gender.M	-0.63** (0.29)	-1.22*** (0.40)	-1.11*** (0.17)	-1.17*** (0.22)	0.48 (0.60)	0.01 (0.28)	0.31 (0.37)
N	2560	2288	2164	4452	2232	1920	4152

Notes: Each data point represents the answer (Join/Not Join) by a Customer in a given round. Standard errors clustered at the session level. Models (1b) and (1c) are restricted to the *Commit* treatment, and models (2b) and (2c) are restricted to the *NoCommit* treatment. We note that the number of observations in models (1b) - (3c) depends on the provider's selected thresholds. Specifically, when providers select a threshold equal to 0 or to Λ , there are only Long Wait or Short Wait signals, respectively.


Figure 25 Low-need Customer Average Joining Probability by Participant (*NoInfo* in Study 1)

Figure 26 Low-need Customer Average Joining Probability Conditional on Signal by Participant (*Commit* in Study 1)




Figure 27 Low-need Customer Average Joining Probability Conditional on Signal by Participant (*NoCommit* in Study 1)

E.1.2 Provider Choices We examine aggregate patterns in providers' threshold choices. Table 5 presents comparisons between average threshold selections and relevant theoretical benchmarks. Table 6 compares average threshold selections between treatments. Table 7 presents comparisons between provider average lying behavior and relevant benchmarks for the case of *NoCommit*. Table 8 presents OLS regressions regarding service providers' lying behavior and social welfare. Figures 28 and 29 show provider threshold distributions at the individual participant level in each signaling treatment.

Table 5 T-Tests: Provider Threshold Selections vs. Theoretical Benchmarks in Study 1

Treatment	Threshold Type	Avg Threshold	Theory/Benchmark	t(df)	p-value	Bootstrap p-value
<i>Commit</i>	Implemented $\theta = \theta'$	2.26	1	$t(7) = 4.36$	0.002	0.009
<i>Commit</i>	Implemented $\theta = \theta'$	2.26	2	$t(7) = 0.89$	0.200	0.230
<i>NoCommit</i>	Implemented θ	2.03	1	$t(6) = 9.25$	< 0.001	< 0.001
<i>NoCommit</i>	Implemented θ	2.03	2	$t(6) = 0.28$	0.391	0.377
<i>NoCommit</i>	Communicated θ'	2.27	1	$t(6) = 5.80$	< 0.001	< 0.001
<i>NoCommit</i>	Communicated θ'	2.27	2	$t(6) = 1.23$	0.131	0.051

Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted values or based on benchmarks of interest. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 6 T-Tests: Treatment Effects, Provider Threshold Selections in Study 1

Comparison	Threshold Type	Avg Threshold	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
<i>Commit</i> vs <i>NoCommit</i>	Implemented vs Implemented	2.26 vs 2.03	$t(13) = 0.69$	0.250	0.256	0.242	0.259
<i>Commit</i> vs <i>NoCommit</i>	Implemented vs Communicated	2.26 vs 2.27	$t(13) = -0.03$	0.486	0.486	0.485	0.463
<i>NoCommit</i>	Communicated vs Implemented (paired)	2.27 vs 2.03	$t(6) = 1.62$	0.077	0.029	–	–

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Paired comparisons used for within-treatment analyses. Bootstrap p-values computed using 9,999 replicates.

Restricting to the second half of the data, the average selected threshold in *Commit* is 2.15. In comparison to the predicted equilibrium threshold of 1: $t(7) = 3.64$, p-value = 0.004, bootstrap t-test p-value = 0.014. For the case of *NoCommit* the average selected *implemented* threshold is 1.85, while the the average *communicated* threshold is 2.21. Comparing these (paired tests): $t(6) = 2.51$, p-value = 0.023, bootstrap t-test p-value = 0.020.

We also conduct a two-sample Kolmogorov-Smirnov test between the distribution of thresholds in *Commit* and the distribution of implemented thresholds in *NoCommit* ($D = 0.21$, p-value < 0.001). And between the distribution of thresholds in *Commit* and the distribution of and communicated thresholds in *NoCommit* ($D = 0.17$, p-value < 0.001).

Table 7 T-Tests: Providers Lying Behavior (*NoCommit*) in Study 1


Sample	Measure	Avg Value	Comparison Benchmark	t(df)	p-value	Bootstrap p-value
All Cases	Lie	0.24	0	$t(6) = 1.62$	0.077	0.023
All Cases	Absolute Lie	0.44	0	$t(6) = 2.40$	0.026	0.010
Lie $\neq 0$ Only	Lie	0.28	0	$t(6) = 0.73$	0.245	0.236
Lie $\neq 0$ Only	Absolute Lie	1.58	0	$t(6) = 8.21$	< 0.001	< 0.001

Notes: The variable *Lie* is defined as $Lie = \theta' - \theta$. We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.


Table 8 OLS Regressions: Service Providers' Lying Behavior in Study 1

	Realized Social Welfare			Implied Social Welfare		
	(1a)	(2a)	(3a)	(1b)	(2b)	(3b)
(Intercept)	625.24*** (3.31)	622.28*** (3.13)	630.48*** (2.21)	613.56*** (6.88)	607.61*** (7.35)	618.98*** (3.92)
<i>Lie</i>	4.52** (1.18)	-	-	8.23*** (0.59)	-	-
$ Lie $	-	0.12 (1.64)	-	-	1.4 (2.59)	-
<i>Lie.Type</i> ($\theta' > \theta$)	-	-	4.21* (2.09)	-	-	12.03** (1.25)
<i>Lie.Type</i> ($\theta' < \theta$)	-	-	-24.56* (7.60)	-	-	-28.75** (3.13)
Round	0.38* (0.12)	0.45* (0.13)	0.38* (0.11)	0.10* (0.03)	0.21* (0.06)	0.12* (0.04)
Gender.M	3.43 (1.82)	4.92* (1.77)	1.22 (1.50)	1.33 (3.53)	4.21 (3.58)	-1.38 (1.79)
N	560	560	560	560	560	560
R^2	0.02	0.01	0.02	0.46	0.08	0.51

Notes: 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.

Figure 28 Provider Threshold Distributions by Participant (*Commit* in Study 1)**Figure 29 Provider Threshold Distributions by Participant (*NoCommit* in Study 1)**

Implemented (θ) — Communicated (θ') —



E.1.3 Social Welfare We examine aggregate social welfare, both realized and implied. Table 9 presents comparisons between average welfare and relevant theoretical benchmarks. Table 10 compares average welfare between treatments. In the *Commit* treatment, we study the implied welfare conditional on some provider thresholds θ' of interest. Table 11 presents comparisons between average implied welfare in *Commit* conditional on θ' and benchmarks/treatments of interest. In the *NoCommit* treatment, we study the implied social welfare conditional on a provider lying behavior: honest ($\theta' = \theta$), overstate ($\theta' > \theta$), and understate ($\theta' < \theta$). Table 12 presents comparisons between average implied welfare in *NoCommit* conditional on lying behavior and treatments of interest. Table 13 presents OLS regressions regarding social welfare. Table 14 presents comparisons between average implied welfare and the *FullInfo* theoretical benchmark.

Table 9 T-Tests: Social Welfare vs. Theoretical Benchmarks in Study 1

Treatment	Welfare Type	Avg Welfare	Theory	t(df)	p-value	Bootstrap p-value
<i>NoInfo</i>	Realized	632.07	580	$t(7) = 15.79$	< 0.001	< 0.001
<i>NoInfo</i>	Implied	615.43	580	$t(7) = 11.99$	< 0.001	< 0.001
<i>Commit</i>	Realized	631.99	634	$t(7) = -0.86$	0.208	0.187
<i>Commit</i>	Implied	613.50	634	$t(7) = -9.85$	< 0.001	< 0.001
<i>NoCommit</i>	Realized	639.91	580	$t(6) = 50.91$	< 0.001	< 0.001
<i>NoCommit</i>	Implied	619.83	580	$t(6) = 28.13$	< 0.001	< 0.001

Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are treatment-specific predictions derived from equilibrium analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 10 T-Tests: Treatment Effects, Social Welfare in Study 1

Comparison	Welfare Type	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Welch Bootstrap p-value
<i>NoInfo</i> vs. <i>Commit</i>	Realized	632.07 vs. 631.99	$t(14) = 0.02$	0.492	0.502	0.492	0.493
<i>NoInfo</i> vs. <i>NoCommit</i>	Realized	632.07 vs. 639.91	$t(13) = -2.11$	0.027	0.023	0.026	0.014
<i>Commit</i> vs. <i>NoCommit</i>	Realized	631.99 vs. 639.91	$t(13) = -2.91$	0.006	0.006	0.006	0.004
<i>NoInfo</i> vs. <i>Commit</i>	Implied	615.43 vs. 613.50	$t(14) = 0.53$	0.301	0.299	0.301	0.304
<i>NoInfo</i> vs. <i>NoCommit</i>	Implied	615.43 vs. 619.83	$t(13) = -1.28$	0.111	0.108	0.104	0.107
<i>Commit</i> vs. <i>NoCommit</i>	Implied	613.50 vs. 619.83	$t(13) = -2.44$	0.015	0.012	0.013	0.024
<i>NoInfo</i> vs. <i>Commit</i>	Implied (2nd Half)	615.78 vs. 614.90	$t(14) = 0.26$	0.399	0.399	0.399	0.404
<i>NoInfo</i> vs. <i>NoCommit</i>	Implied (2nd Half)	615.78 vs. 621.73	$t(13) = -2.03$	0.032	0.032	0.029	0.032
<i>Commit</i> vs. <i>NoCommit</i>	Implied (2nd Half)	614.90 vs. 621.73	$t(13) = -2.39$	0.016	0.032	0.015	0.033

Notes: We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates. “2nd Half” refers to observations from the second half of each session.

Table 11 T-Tests: Treatment Effects, Implied Social Welfare by Thresholds θ' of Interest (*Commit*) in Study 1

Comparison	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
<i>Commit</i> ($\theta' = 0$) vs. <i>NoInfo</i>	611.38 vs. 615.43	$t(11) = -0.70$	0.248	0.241	0.273	0.266
<i>Commit</i> ($\theta' = 1$) vs. <i>NoInfo</i>	621.19 vs. 615.43	$t(10) = 1.11$	0.145	0.147	0.156	0.160
<i>Commit</i> ($\theta' = 4$) vs. <i>NoInfo</i>	603.83 vs. 615.43	$t(13) = -2.52$	0.013	0.014	0.014	0.014
<i>Commit</i> ($\theta' = 1$) vs. <i>Commit</i> Theory	621.19 vs. 634	$t(3) = -2.98$	0.029	0.023	—	—

Notes: We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 12 T-Tests: Treatment Effects, Implied Social Welfare by Lying Behavior (*NoCommit*) in Study 1

Comparison	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
<i>NoCommit</i> (<i>Honest</i>) vs. <i>NoInfo</i>	618.38 vs. 615.43	$t(13) = 0.90$	0.191	0.190	0.182	0.177
<i>NoCommit</i> (<i>Honest</i>) vs. <i>Commit</i>	618.38 vs. 613.50	$t(13) = 2.06$	0.029	0.029	0.028	0.036
<i>NoCommit</i> (<i>Understate</i>) vs. <i>NoInfo</i>	594.59 vs. 615.43	$t(13) = -3.86$	0.001	0.002	0.002	0.005
<i>NoCommit</i> (<i>Understate</i>) vs. <i>Commit</i>	594.59 vs. 613.50	$t(13) = -3.86$	0.001	0.002	0.003	0.007
<i>NoCommit</i> (<i>Overstate</i>) vs. <i>NoInfo</i>	629.10 vs. 615.43	$t(12) = 2.84$	0.007	0.008	0.010	0.015
<i>NoCommit</i> (<i>Overstate</i>) vs. <i>Commit</i>	629.10 vs. 613.50	$t(12) = 3.77$	0.001	0.002	0.004	0.008
<i>NoCommit</i> (<i>Overstate</i>) vs. <i>Commit</i> Theory	629.10 vs. 634	$t(5) = -1.24$	0.113	0.069	–	–

Notes: *Honest* ($\theta' = \theta$), *overstate* ($\theta' > \theta$), and *understate* ($\theta' < \theta$). We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 13 OLS Regressions: Social Welfare in Study 1

	Realized Social Welfare					Implied Social Welfare				
	(1a)	(2a)	(3a) [†]	(4a)	(5a) [‡]	(1b)	(2b)	(3b) [†]	(4b)	(5b) [‡]
(Intercept)	630.01*** (4.39)	630.01*** (4.40)	630.74*** (2.13)	630.01*** (4.35)	630.01*** (4.35)	614.58*** (4.02)	614.58*** (4.03)	615.14*** (1.83)	614.58*** (3.99)	614.58*** (3.99)
<i>Commit</i>	-5.47 (4.71)	- (2.73)	-6.20* (2.73)	-5.47 (4.66)	-5.47 (4.66)	-4.11 (4.42)	- (2.59)	-4.68* (4.37)	-4.11 (4.38)	-4.11 (4.38)
<i>NoCommit</i>	- (4.71)	0.73 (4.89)	- (4.83)	0.73 (4.83)	- (4.83)	- (4.43)	0.56 (4.43)	- (4.38)	0.56 (4.38)	- (4.38)
Round	0.10 (0.10)	0.10 (0.10)	0.45** (0.11)	0.10 (0.10)	0.10 (0.10)	0.04 (0.07)	0.04 (0.07)	0.23** (0.06)	0.04 (0.07)	0.04 (0.07)
<i>Commit</i> *Round	0.26* (0.13)	- (0.15)	-0.08 (0.14)	0.26* (0.13)	0.26* (0.13)	0.11 (0.09)	- (0.09)	-0.08 (0.08)	0.11 (0.09)	0.11 (0.09)
<i>NoCommit</i> *Round	- (0.15)	0.35* (0.15)	- (0.14)	0.35* (0.14)	- (0.14)	- (0.09)	0.19* (0.09)	- (0.09)	0.19* (0.09)	- (0.09)
<i>NoCommit</i> (<i>Honest</i>)	-	-	-	-	1.06 (4.97)	-	-	-	-	2.34 (4.14)
<i>NoCommit</i> (<i>Overstate</i>)	-	-	-	-	10.15* (4.61)	-	-	-	-	10.38* (4.80)
<i>NoCommit</i> (<i>Understate</i>)	-	-	-	-	-9.82 (17.60)	-	-	-	-	-21.36*** (5.47)
<i>NoCommit</i> (<i>Honest</i>)*Round	-	-	-	-	0.36* (0.15)	-	-	-	-	0.06 (0.08)
<i>NoCommit</i> (<i>Overstate</i>)*Round	-	-	-	-	0.15 (0.11)	-	-	-	-	0.22 (0.14)
<i>NoCommit</i> (<i>Understate</i>)*Round	-	-	-	-	-0.50 (0.94)	-	-	-	-	-0.21 (0.16)
N	1280	1200	1200	1840	1840	1280	1200	1200	1840	1840
R ²	0.00	0.01	0.01	0.01	0.01	0.01	0.04	0.09	0.05	0.16

Notes: 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 *NoInfo* treatment, and take as baseline the *NoCommit* treatment to compare it with the *Commit* treatment.

‡: Models (5a) and (5b) consider the disaggregated *NoCommit* treatment: *honest* ($\theta' = \theta$), *overstate* ($\theta' > \theta$), and *understate* ($\theta' < \theta$).

Table 14 T-Tests: Implied Social Welfare vs. *FullInfo* Theoretical Benchmark in Study 1

Treatment	Avg Welfare	Theory	t(df)	p-value	Bootstrap p-value
<i>NoInfo</i>	615.43	620.5	$t(7) = -1.71$	0.065	0.061
<i>Commit</i>	613.50	620.5	$t(7) = -3.36$	0.006	0.010
<i>Commit</i> ($\theta = 1$)	621.19	620.5	$t(3) = 0.16$	0.441	0.404
<i>NoCommit</i>	619.83	620.5	$t(6) = -0.47$	0.327	0.280
<i>NoCommit</i> (<i>Honest</i>)	618.38	620.5	$t(6) = -2.59$	0.020	0.004
<i>NoCommit</i> (<i>Overstate</i>)	629.10	620.5	$t(5) = 2.19$	0.034	0.170
<i>NoCommit</i> (<i>Understate</i>)	594.59	620.5	$t(6) = -5.53$	< 0.001	0.003

Notes: We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

E.1.4 Practice Rounds We examine aggregate patterns in customers' joining behavior and implied welfare. Table 15 presents comparisons between average customer joining probabilities and relevant benchmarks. Table 16 presents comparisons of average implied welfare between practice rounds and actual treatments. Figure 30 shows plots of customer average joining behavior and implied welfare.

Table 15 T-Tests: Customer Joining Behavior vs. Theoretical Benchmarks in Study 1 Practice Rounds

Customer Type	Queue Length	Avg Join Prob.	Theory	t(df)	p-value	Bootstrap p-value
High-need	0	0.960	1	$t(22) = -4.69$	< 0.001	< 0.001
High-need	1	0.974	1	$t(22) = -4.21$	< 0.001	< 0.001
High-need	2	0.965	1	$t(22) = -4.27$	< 0.001	< 0.001
High-need	3	0.868	1	$t(22) = -7.58$	< 0.001	< 0.001
Low-need	0	0.949	1	$t(22) = -6.23$	< 0.001	< 0.001
Low-need	1	0.919	1	$t(22) = -6.68$	< 0.001	< 0.001
Low-need	2	0.168	0	$t(22) = 9.03$	< 0.001	< 0.001
Low-need	3	0.077	0	$t(22) = 5.84$	< 0.001	< 0.001


Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted joining probabilities for each customer type, or based on benchmarks of interest. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 16 T-Tests: Implied Social Welfare in Main Experiment vs. Practice Rounds in Study 1

Treatment	Avg Welfare (Main vs. Practice)	t(df)	p-value	Bootstrap p-value
<i>NoInfo</i> vs. <i>NoInfo</i> (Practice Rounds)	615.43 vs. 610.15	$t(7) = -3.09$	0.008	0.022
<i>Commit</i> vs. <i>Commit</i> (Practice Rounds)	613.50 vs. 607.02	$t(7) = -4.45$	0.001	< 0.001
<i>NoCommit</i> vs. <i>NoCommit</i> (Practice Rounds)	619.83 vs. 615.87	$t(6) = -1.65$	0.074	0.095
All Treatments(Practice Rounds) vs. <i>FullInfo</i> Theory	610.80 vs. 620.5	$t(22) = -8.43$	< 0.001	< 0.001

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates. Paired t-tests compare implied social welfare between main and practice rounds.

Figure 30 Customer Joining Behavior and Implied Social Welfare (Practice Rounds) in Study 1



E.2 Study 2


E.2.1 Customer Choices We examine aggregate patterns in customers' joining behavior. Table 17 presents comparisons between average customer joining probabilities and relevant theoretical benchmarks and treatments. Finally, Figure 31 show average (across rounds) joining probabilities for low-need customers at the individual participant level.

Table 17 T-Tests: Customer Joining Behavior, Comparisons of Interest in Study 2

Type	Comparison	Signal	Threshold θ'	Avg Prob	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
High-need	<i>CommitBOT</i> vs <i>CommitBOT</i> Theory	Short Wait	1	1.000 vs 1	—	—	—	—	—
High-need	<i>CommitBOT</i> vs <i>CommitBOT</i> Theory	Long Wait	1	0.997 vs 1	$t(7) = -3.86$	0.003	0.024	—	—
Low-need	<i>CommitBOT</i> vs <i>CommitBOT</i> Theory	Short Wait	1	0.977 vs 1	$t(7) = -1.43$	0.098	0.035	—	—
Low-need	<i>CommitBOT</i> vs <i>CommitBOT</i> Theory	Long Wait	1	0.171 vs 0	$t(7) = 4.67$	0.001	0.004	—	—
Low-need	<i>CommitBOT</i> vs <i>Commit</i> ($\theta' = 1$)	Short Wait	1	0.977 vs 0.937	$t(10) = 0.83$	0.212	NA	0.289	0.242
Low-need	<i>CommitBOT</i> vs <i>NoCommit</i> ($\theta' = 1$)	Short Wait	1	0.977 vs 0.975	$t(11) = 0.08$	0.467	NA	0.469	0.432
Low-need	<i>CommitBOT</i> vs <i>Commit</i> ($\theta' = 1$)	Long Wait	1	0.171 vs 0.220	$t(10) = -0.82$	0.216	0.211	0.207	0.192
Low-need	<i>CommitBOT</i> vs <i>NoCommit</i> ($\theta' = 1$)	Long Wait	1	0.171 vs 0.191	$t(11) = -0.37$	0.357	0.357	0.349	0.353

Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted joining probabilities for each customer type. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates. Constant values resulted in omitted t-stats.

Figure 31 Low-need Customer Average Joining Probability Conditional on Signal by Participant (*CommitBOT* in Study 2)



E.2.2 Social Welfare We examine aggregate social welfare, both realized and implied. Table 18 presents comparisons between average welfare and relevant theoretical benchmarks and treatments. Table 19 presents OLS regressions regarding social welfare.

Table 18 T-Tests: Social Welfare, Comparisons of Interest in Study 2

Comparison	Welfare Type	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Welch Bootstrap p-value
<i>CommitBOT</i> vs. <i>CommitBOT</i> Theory	Realized	643.71 vs. 634	$t(7) = 4.14$	0.002	0.011	—	—
<i>CommitBOT</i> vs. <i>CommitBOT</i> Theory	Implied	626.26 vs. 634	$t(7) = -4.59$	0.001	0.005	—	—
<i>CommitBOT</i> vs. <i>FullInfo</i> Theory	Implied	626.26 vs. 620.5	$t(7) = 3.42$	0.005	0.006	—	—
<i>CommitBOT</i> vs. <i>NoInfo</i>	Realized	643.71 vs. 632.07	$t(14) = 2.87$	0.006	0.007	0.007	0.004
<i>CommitBOT</i> vs. <i>Commit</i>	Realized	643.71 vs. 631.99	$t(14) = 3.55$	0.001	0.002	0.001	0.002
<i>CommitBOT</i> vs. <i>NoCommit</i>	Realized	643.71 vs. 639.91	$t(13) = 1.38$	0.094	0.096	0.089	0.104
<i>CommitBOT</i> vs. <i>NoCommit(Honest)</i>	Realized	643.71 vs. 639.09	$t(13) = 1.61$	0.065	0.064	0.061	0.077
<i>CommitBOT</i> vs. <i>NoCommit(Overstate)</i>	Realized	643.71 vs. 650.57	$t(12) = -1.28$	0.113	0.113	0.142	0.097
<i>CommitBOT</i> vs. <i>NoCommit(Understate)</i>	Realized	643.71 vs. 602.06	$t(13) = 3.53$	0.002	0.001	0.007	0.005
<i>CommitBOT</i> vs. <i>NoInfo</i>	Implied	626.26 vs. 615.43	$t(14) = 3.18$	0.003	0.003	0.004	0.004
<i>CommitBOT</i> vs. <i>Commit</i>	Implied	626.26 vs. 613.50	$t(14) = 4.76$	< 0.001	< 0.001	< 0.001	< 0.001
<i>CommitBOT</i> vs. <i>Commit</i> ($\theta = 1$)	Implied	626.26 vs. 621.19	$t(10) = 1.34$	0.104	0.110	0.166	0.160
<i>CommitBOT</i> vs. <i>NoCommit</i>	Implied	626.26 vs. 619.83	$t(13) = 2.87$	0.006	0.007	0.006	0.003
<i>CommitBOT</i> vs. <i>NoCommit(Honest)</i>	Implied	626.26 vs. 618.38	$t(13) = 4.01$	< 0.001	< 0.001	< 0.001	0.001
<i>CommitBOT</i> vs. <i>NoCommit(Overstate)</i>	Implied	626.26 vs. 629.10	$t(12) = -0.73$	0.239	0.244	0.263	0.320
<i>CommitBOT</i> vs. <i>NoCommit(Understate)</i>	Implied	626.26 vs. 594.59	$t(13) = 6.71$	< 0.001	< 0.001	< 0.001	0.001

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 19 OLS Regressions: Social Welfare in Study 2

	Realized Social Welfare				Implied Social Welfare			
	(1a)	(2a) [†]	(3a) [‡]	(4a)	(1b)	(2b) [†]	(3b) [‡]	(4b)
(Intercept)	630.01*** (4.39)	624.54*** (1.70)	630.74*** (2.13)	630.01*** (4.35)	614.58*** (4.02)	610.46*** (1.82)	615.14*** (1.83)	614.58*** (3.96)
<i>CommitBOT</i>	10.69* (4.98)	16.16*** (2.91)	9.96** (3.19)	10.69* (4.91)	9.91* (4.40)	14.02*** (2.54)	9.35** (2.55)	9.91* (4.33)
Round	0.10 (0.10)	0.36** (0.09)	0.45** (0.11)	0.10 (0.09)	0.04 (0.07)	0.15** (0.06)	0.23** (0.06)	0.04 (0.07)
<i>CommitBOT</i> *Round	0.05 (0.10)	-0.22** (0.10)	-0.30* (0.12)	0.05 (0.10)	0.04 (0.09)	-0.06 (0.08)	-0.14* (0.08)	0.04 (0.09)
<i>Commit</i>	- (4.63)	- (4.63)	- (4.63)	-5.47 (4.63)	- (4.63)	- (4.63)	- (4.63)	-4.11 (4.35)
<i>NoCommit</i>	- (4.80)	- (4.80)	- (4.80)	0.73 (4.80)	- (4.80)	- (4.80)	- (4.80)	0.56 (4.35)
<i>Commit</i> *Round	- (0.13)	- (0.13)	- (0.13)	0.26* (0.13)	- (0.13)	- (0.13)	- (0.13)	0.11 (0.09)
<i>NoCommit</i> *Round	- (0.14)	- (0.14)	- (0.14)	0.35* (0.14)	- (0.14)	- (0.14)	- (0.14)	0.19* (0.09)
N	1280	1280	1200	2480	1280	1280	1200	2480
R ²	0.01	0.01	0.01	0.01	0.15	0.28	0.12	0.15

Notes: 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 (2a) and (2b) do not consider the *NoInfo* treatment, and take as baseline the *Commit* treatment to compare it with the *CommitBOT* treatment.

[‡]: Models (3a) and (3b) do not consider the *NoInfo* treatment, and take as baseline the *NoCommit* treatment to compare it with the *CommitBOT* treatment.

E.2.3 Practice Rounds We examine aggregate patterns in customers' joining behavior and implied welfare. Table 20 presents comparisons between average customer joining probabilities and relevant benchmarks. Table 21 presents comparisons of average implied welfare between practice rounds and actual treatments. Figure 32 shows plots of customer average joining behavior and implied welfare.

Table 20 T-Tests: Customer Joining Behavior vs. Theoretical Benchmarks in Study 2 Practice Rounds

Type	Queue Length	Avg Join Prob.	Theory	t(df)	p-value	Bootstrap p-value
High-need	0	0.981	1	$t(7) = -3.00$	0.010	0.004
High-need	1	0.976	1	$t(7) = -2.70$	0.015	0.003
High-need	2	0.968	1	$t(7) = -2.82$	0.013	0.008
High-need	3	0.937	1	$t(7) = -3.44$	0.005	0.006
Low-need	0	0.958	1	$t(7) = -2.89$	0.011	0.006
Low-need	1	0.953	1	$t(7) = -3.98$	0.003	0.006
Low-need	2	0.122	0	$t(7) = 2.63$	0.017	0.003
Low-need	3	0.053	0	$t(7) = 3.16$	0.008	0.003


Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted joining probabilities for each customer type, or based on benchmarks of interest. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates

Table 21 T-Tests: Implied Social Welfare in Main Experiment vs. Practice Rounds in Study 2

Treatment	Avg Welfare (Main vs. Practice)	t(df)	p-value	Bootstrap p-value
<i>CommitBOT vs. CommitBOT (Practice Rounds)</i>	626.26 vs. 613.76	$t(7) = -8.71$	< 0.001	< 0.001
<i>CommitBOT (Practice Rounds) vs. FullInfo Theory</i>	613.76 vs. 620.50	$t(7) = -3.98$	0.002	< 0.001

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set based on the expected direction of the treatment effect. Bootstrap p-values computed using 9,999 replicates. Paired t-tests compare implied social welfare between main and practice rounds.

Figure 32 Customer Joining Behavior and Implied Social Welfare (Practice Rounds) in Study 2



E.3 Study 3

E.3.1 Customer Choices We examine aggregate patterns in customers' joining behavior. Table 22 presents comparisons between average low-need customer joining probabilities and relevant theoretical benchmarks. Table 23 compares average joining probabilities for low-need customers between treatments. Table 24 presents logistic regression results for low-need customers' joining probabilities across treatments. Finally, Figures 33–35 show average (across rounds) joining probabilities for low-need customers at the individual participant level in each treatment.

Table 22 T-Tests: Low-Need Customer Joining Behavior vs. Theoretical Benchmarks in Study 3

Treatment	Signal	Threshold θ^*	Avg Prob	Theory/Benchmark	t(df)	p-value	Bootstrap p-value
<i>NoInfo</i>	–	–	0.571	1	$t(2) = -8.50$	0.007	0.076
<i>NoInfo</i>	–	–	0.571	0.5	$t(2) = 1.40$	0.148	0.036
<i>NoInfo</i>	–	–	0.571	0	$t(2) = 11.31$	0.004	0.036
<i>Commit</i>	Short Wait	All	0.880	1	$t(4) = -5.17$	0.003	0.003
<i>Commit</i>	Long Wait	All	0.238	0	$t(4) = 6.08$	0.002	< 0.001
<i>Commit</i>	Short Wait	1	0.910	1	$t(4) = -3.83$	0.009	0.106
<i>Commit</i>	Short Wait	2	0.913	1	$t(4) = -5.15$	0.003	0.016
<i>Commit</i>	Short Wait	3	0.910	1	$t(4) = -4.33$	0.006	0.010
<i>Commit</i>	Short Wait	4	0.894	1	$t(4) = -4.22$	0.006	0.064
<i>Commit</i>	Short Wait	5	0.849	1	$t(3) = -2.80$	0.034	0.004
<i>Commit</i>	Short Wait	6	0.702	1	$t(3) = -3.74$	0.017	0.005
<i>Commit</i>	Short Wait	7	0.458	1	$t(2) = -4.91$	0.019	0.036
<i>Commit</i>	Short Wait	8	0.555	1	$t(4) = -8.93$	< 0.001	< 0.001
<i>Commit</i>	Long Wait	0	0.481	1	$t(4) = -11.49$	< 0.001	< 0.001
<i>Commit</i>	Long Wait	1	0.427	0.65	$t(4) = -3.16$	0.017	0.025
<i>Commit</i>	Long Wait	2	0.285	0	$t(4) = 10.22$	< 0.001	0.008
<i>Commit</i>	Long Wait	3	0.180	0	$t(4) = 4.83$	0.004	< 0.001
<i>Commit</i>	Long Wait	4	0.128	0	$t(4) = 2.32$	0.040	0.014
<i>Commit</i>	Long Wait	5	0.088	0	$t(3) = 1.92$	0.074	0.046
<i>Commit</i>	Long Wait	6	0.057	0	$t(3) = 2.12$	0.062	0.068
<i>Commit</i>	Long Wait	7	0.208	0	$t(2) = 1.88$	0.010	0.296
<i>NoCommit</i>	Short Wait	All	0.837	1	$t(6) = -8.20$	< 0.001	< 0.001
<i>NoCommit</i>	Long Wait	All	0.308	1	$t(6) = -21.16$	< 0.001	< 0.001
<i>NoCommit</i>	Short Wait	1	0.841	1	$t(6) = -4.93$	0.001	0.005
<i>NoCommit</i>	Short Wait	2	0.828	1	$t(6) = -5.58$	< 0.001	< 0.001
<i>NoCommit</i>	Short Wait	3	0.875	1	$t(6) = -3.79$	0.004	0.001
<i>NoCommit</i>	Short Wait	4	0.856	1	$t(6) = -6.45$	< 0.001	< 0.001
<i>NoCommit</i>	Short Wait	5	0.803	1	$t(6) = -9.00$	< 0.001	< 0.001
<i>NoCommit</i>	Short Wait	6	0.779	1	$t(6) = -9.73$	< 0.001	< 0.001
<i>NoCommit</i>	Short Wait	7	0.718	1	$t(5) = -9.00$	< 0.001	0.001
<i>NoCommit</i>	Short Wait	8	0.722	1	$t(5) = -5.65$	0.001	0.071
<i>NoCommit</i>	Long Wait	0	0.403	1	$t(6) = -21.15$	< 0.001	< 0.001
<i>NoCommit</i>	Long Wait	1	0.393	1	$t(6) = -26.54$	< 0.001	< 0.001
<i>NoCommit</i>	Long Wait	2	0.325	1	$t(6) = -27.22$	< 0.001	< 0.001
<i>NoCommit</i>	Long Wait	3	0.308	1	$t(6) = -14.98$	< 0.001	0.003
<i>NoCommit</i>	Long Wait	4	0.230	1	$t(6) = -19.22$	< 0.001	< 0.001
<i>NoCommit</i>	Long Wait	5	0.223	1	$t(6) = -23.81$	< 0.001	< 0.001
<i>NoCommit</i>	Long Wait	6	0.249	1	$t(6) = -10.44$	< 0.001	0.002
<i>NoCommit</i>	Long Wait	7	0.223	1	$t(5) = -16.89$	< 0.001	< 0.001

Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted joining probabilities, or based on benchmarks of interest. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 23 T-Tests: Treatment Effects, Low-Need Customer Joining Behavior in Study 3

Comparison	Signal	Threshold	Avg Prob	t(df)	p-value	Bootstrap p-value	Welch p-value	Welch Bootstrap p-value
<i>Commit vs NoInfo</i>	Short Wait	All	0.880 vs 0.571	$t(6) = 6.42$	< 0.001	0.004	0.006	0.016
<i>Commit vs NoInfo</i>	Long Wait	All	0.238 vs 0.571	$t(6) = -5.20$	0.001	0.005	0.006	0.016
<i>Commit vs NoInfo</i>	Short Wait	1	0.910 vs 0.571	$t(6) = 7.03$	< 0.001	0.003	0.005	0.008
<i>Commit vs NoInfo</i>	Short Wait	2	0.913 vs 0.571	$t(6) = 7.92$	< 0.001	0.003	0.006	0.008
<i>Commit vs NoInfo</i>	Short Wait	3	0.910 vs 0.571	$t(6) = 7.37$	< 0.001	0.003	0.006	0.023
<i>Commit vs NoInfo</i>	Short Wait	4	0.894 vs 0.571	$t(6) = 6.52$	< 0.001	0.003	0.005	0.003
<i>Commit vs NoInfo</i>	Short Wait	5	0.849 vs 0.571	$t(5) = 3.67$	0.007	NA	0.006	0.093
<i>Commit vs NoInfo</i>	Short Wait	6	0.702 vs 0.571	$t(5) = 1.27$	0.130	0.124	0.113	0.182
<i>Commit vs NoInfo</i>	Short Wait	7	0.458 vs 0.571	$t(4) = -0.93$	0.203	NA	0.213	0.174
<i>Commit vs NoInfo</i>	Short Wait	8	0.555 vs 0.571	$t(6) = -0.21$	0.421	0.423	0.416	0.405
<i>Commit vs NoInfo</i>	Long Wait	0	0.481 vs 0.571	$t(6) = -1.26$	0.126	0.126	0.122	0.079
<i>Commit vs NoInfo</i>	Long Wait	1	0.427 vs 0.571	$t(6) = -1.42$	0.102	0.096	0.074	0.063
<i>Commit vs NoInfo</i>	Long Wait	2	0.285 vs 0.571	$t(6) = -5.44$	< 0.001	0.004	0.006	0.015
<i>Commit vs NoInfo</i>	Long Wait	3	0.180 vs 0.571	$t(6) = -6.29$	< 0.001	0.003	0.001	0.010
<i>Commit vs NoInfo</i>	Long Wait	4	0.128 vs 0.571	$t(6) = -5.35$	< 0.001	0.005	< 0.001	0.014
<i>Commit vs NoInfo</i>	Long Wait	5	0.088 vs 0.571	$t(5) = -6.98$	< 0.001	0.007	< 0.001	0.003
<i>Commit vs NoInfo</i>	Long Wait	6	0.057 vs 0.571	$t(5) = -9.96$	< 0.001	0.006	0.001	0.003
<i>Commit vs NoInfo</i>	Long Wait	7	0.208 vs 0.571	$t(4) = -2.99$	0.020	NA	0.032	0.038
<i>NoCommit vs NoInfo</i>	Short Wait	All	0.837 vs 0.571	$t(8) = 6.09$	< 0.001	0.001	0.011	0.023
<i>NoCommit vs NoInfo</i>	Long Wait	All	0.308 vs 0.571	$t(8) = -4.39$	0.001	0.004	0.006	0.012
<i>NoCommit vs NoInfo</i>	Short Wait	1	0.841 vs 0.571	$t(8) = 4.56$	< 0.001	0.003	0.006	0.005
<i>NoCommit vs NoInfo</i>	Short Wait	2	0.828 vs 0.571	$t(8) = 4.49$	0.001	0.004	0.007	0.051
<i>NoCommit vs NoInfo</i>	Short Wait	3	0.875 vs 0.571	$t(8) = 5.06$	< 0.001	0.002	0.004	0.015
<i>NoCommit vs NoInfo</i>	Short Wait	4	0.856 vs 0.571	$t(8) = 6.18$	< 0.001	0.002	0.008	0.023
<i>NoCommit vs NoInfo</i>	Short Wait	5	0.803 vs 0.571	$t(8) = 5.09$	< 0.001	0.002	0.014	0.036
<i>NoCommit vs NoInfo</i>	Short Wait	6	0.779 vs 0.571	$t(8) = 4.46$	0.001	0.004	0.018	0.036
<i>NoCommit vs NoInfo</i>	Short Wait	7	0.718 vs 0.571	$t(7) = 2.62$	0.017	0.020	0.037	0.075
<i>NoCommit vs NoInfo</i>	Short Wait	8	0.722 vs 0.571	$t(7) = 1.93$	0.048	0.036	0.039	0.029
<i>NoCommit vs NoInfo</i>	Long Wait	0	0.403 vs 0.571	$t(8) = -3.11$	0.007	0.008	0.027	0.036
<i>NoCommit vs NoInfo</i>	Long Wait	1	0.393 vs 0.571	$t(8) = -3.77$	0.003	0.004	0.016	0.021
<i>NoCommit vs NoInfo</i>	Long Wait	2	0.325 vs 0.571	$t(8) = -4.96$	< 0.001	0.001	0.003	0.006
<i>NoCommit vs NoInfo</i>	Long Wait	3	0.308 vs 0.571	$t(8) = -3.32$	0.005	0.010	0.005	0.015
<i>NoCommit vs NoInfo</i>	Long Wait	4	0.230 vs 0.571	$t(8) = -4.85$	< 0.001	0.002	0.002	0.006
<i>NoCommit vs NoInfo</i>	Long Wait	5	0.223 vs 0.571	$t(8) = -5.81$	< 0.001	0.002	0.003	0.006
<i>NoCommit vs NoInfo</i>	Long Wait	6	0.249 vs 0.571	$t(8) = -2.74$	0.013	0.013	0.003	0.006
<i>NoCommit vs NoInfo</i>	Long Wait	7	0.223 vs 0.571	$t(7) = -4.62$	0.001	0.003	0.002	0.002
<i>Commit vs NoCommit</i>	Short Wait	All	0.880 vs 0.837	$t(10) = 1.41$	0.093	0.090	0.095	0.103
<i>Commit vs NoCommit</i>	Long Wait	All	0.238 vs 0.308	$t(10) = -1.38$	0.099	0.097	0.101	0.128
<i>Commit vs NoCommit</i>	Short Wait	1	0.910 vs 0.841	$t(10) = 1.60$	0.070	0.072	0.056	0.051
<i>Commit vs NoCommit</i>	Short Wait	2	0.913 vs 0.828	$t(10) = 2.14$	0.029	0.017	0.019	0.005
<i>Commit vs NoCommit</i>	Short Wait	3	0.910 vs 0.875	$t(10) = 0.812$	0.218	0.230	0.195	0.168
<i>Commit vs NoCommit</i>	Short Wait	4	0.894 vs 0.856	$t(10) = 1.122$	0.144	0.148	0.143	0.118
<i>Commit vs NoCommit</i>	Short Wait	5	0.849 vs 0.803	$t(9) = 0.94$	0.184	0.188	0.234	0.322
<i>Commit vs NoCommit</i>	Short Wait	6	0.702 vs 0.779	$t(9) = -1.18$	0.133	0.146	0.205	0.137
<i>Commit vs NoCommit</i>	Short Wait	7	0.458 vs 0.718	$t(7) = -3.05$	0.009	NA	0.067	0.036
<i>Commit vs NoCommit</i>	Short Wait	8	0.555 vs 0.722	$t(9) = -2.39$	0.020	0.019	0.020	0.012
<i>Commit vs NoCommit</i>	Long Wait	0	0.481 vs 0.403	$t(10) = 1.55$	0.076	0.080	0.093	0.172
<i>Commit vs NoCommit</i>	Long Wait	1	0.427 vs 0.393	$t(10) = 0.52$	0.306	0.310	0.334	0.322
<i>Commit vs NoCommit</i>	Long Wait	2	0.285 vs 0.325	$t(10) = -1.06$	0.157	0.163	0.156	0.135
<i>Commit vs NoCommit</i>	Long Wait	3	0.180 vs 0.308	$t(10) = -2.01$	0.036	0.029	0.029	0.017
<i>Commit vs NoCommit</i>	Long Wait	4	0.128 vs 0.230	$t(10) = -1.54$	0.078	0.070	0.087	0.133
<i>Commit vs NoCommit</i>	Long Wait	5	0.088 vs 0.223	$t(9) = -2.44$	0.018	0.019	0.027	0.037
<i>Commit vs NoCommit</i>	Long Wait	6	0.057 vs 0.249	$t(9) = -1.92$	0.043	0.029	0.020	0.016
<i>Commit vs NoCommit</i>	Long Wait	7	0.208 vs 0.223	$t(7) = -0.16$	0.439	NA	0.452	0.431

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 24 Logistic Regressions: Customer Strategies in *NoInfo*, *Commit*, and *NoCommit* in Study 3

	<i>NoInfo</i> $\mathbb{P}(\text{Join})$ (1a)	<i>Commit & NoCommit</i> $\mathbb{P}(\text{Join} \text{Short Wait})$ (1b) (2b) (3b)			<i>Commit & NoCommit</i> $\mathbb{P}(\text{Join} \text{Long Wait})$ (1c) (2c) (3c)		
(Intercept)	0.07 (0.51)	5.17*** (0.75)	2.20** (0.63)	4.07*** (0.58)	0.35 (0.27)	-0.06 (0.40)	0.27 (0.38)
<i>NoCommit</i>	-	-	-	-1.29*** (0.27)	-	-	-0.29 (0.26)
<i>Communicated Threshold</i>	-	-0.44*** (0.06)	-0.14*** (0.02)	-0.42*** (0.05)	-0.49*** (0.08)	-0.19*** (0.04)	-0.49*** (0.07)
<i>NoCommit * Communicated Threshold</i>	-	-	-	0.28*** (0.05)	-	-	0.29*** (0.08)
Round	-0.01*** (0.00)	0.00 (0.00)	0.01** (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)
Gender.M	0.28 (0.35)	-0.88** (0.34)	-0.11 (0.38)	-0.39 (0.29)	-0.14 (0.16)	-0.14 (0.30)	-0.14 (0.20)
N	1920	2896	3704	6600	3096	4384	7480

Notes: Each data point represents the answer (Join/Not Join) by a Customer in a given round. Standard errors clustered at the session level. We note that the number of observations in models (1b) - (3c) depends on the provider's selected thresholds. Specifically, when providers select a threshold equal to 0 or to Λ , there are only Long Wait or Short Wait signals, respectively. Models (1b) and (1c) are restricted to the *Commit* treatment, and models (2b) and (2c) are restricted to the *NoCommit* treatment.

Figure 33 Low-need Customer Average Joining Probability Conditional on Signal by Participant (*NoInfo* in Study 3)






Figure 34 Low-need Customer Average Joining Probability Conditional on Signal by Participant (*Commit* in Study 3)

Figure 35 Low-need Customer Average Joining Probability Conditional on Signal by Participant (*NoCommit* in Study 3)



E.3.2 Provider Choices We examine aggregate patterns in providers' threshold choices. Table 25 presents comparisons between average threshold selections and relevant theoretical benchmarks. Table 26 compares average threshold selections between treatments. Table 27 presents comparisons between provider average lying behavior and relevant benchmarks for the case of *NoCommit*. Table 28 presents OLS regressions regarding service providers' lying behavior and social welfare. Figures 36 and 37 show provider threshold distributions at the individual participant level in each signaling treatment.

Table 25 T-Tests: Provider Threshold Selections vs. Theoretical Benchmarks in Study 3

Treatment	Threshold Type	Avg Threshold	Theory/Benchmark	t(df)	p-value	Bootstrap p-value
<i>Commit</i>	Implemented	2.98	2	$t(4) = 3.99$	0.008	0.069
<i>Commit</i>	Implemented	2.98	4	$t(4) = -4.20$	0.006	< 0.001
<i>NoCommit</i>	Implemented	2.46	2	$t(6) = 1.11$	0.154	0.169
<i>NoCommit</i>	Implemented	2.46	4	$t(6) = -3.72$	0.005	0.002
<i>NoCommit</i>	Communicated	3.16	2	$t(6) = 3.81$	0.004	0.250
<i>NoCommit</i>	Communicated	3.16	4	$t(6) = -2.74$	0.017	0.001

Notes: We use session-level averages as the unit of analysis. Theoretical benchmarks are based on predicted values or based on benchmarks of interest. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 26 T-Tests: Treatment Effects, Provider Threshold Selections in Study 3

Comparison	Threshold Type	Avg Threshold	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
<i>Commit vs NoCommit</i>	Implemented vs Implemented	2.98 vs 2.46	$t(10) = 0.96$	0.180	0.184	0.155	0.147
<i>Commit vs NoCommit</i>	Implemented vs Communicated	2.98 vs 3.16	$t(10) = -0.45$	0.330	0.336	0.319	0.365
<i>NoCommit</i>	Communicated vs Implemented (paired)	3.16 vs 2.46	$t(6) = 2.69$	0.018	0.017	—	—

Notes: We use session-level averages as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Paired comparisons used for within-treatment analyses. Bootstrap p-values computed using 9,999 replicates.

Restricting to the second half of the data, the average selected threshold in *Commit* is 2.83. In comparison to the predicted equilibrium threshold of 2: $t(4) = 5.00$, p-value = 0.004, bootstrap t-test p-value = 0.074. For the case of *NoCommit* the average selected *implemented* threshold is 2.21, while the the average *communicated* threshold is 3.12. Comparing these (paired tests): $t(6) = 2.86$, p-value = 0.014, bootstrap t-test p-value = 0.004.

We also conduct a two-sample Kolmogorov-Smirnov test between the distribution of thresholds in *Commit* and the distribution of implemented thresholds in *NoCommit* ($D = 0.15$, p-value < 0.001). Also between the distribution of thresholds in *Commit* and the distribution of and communicated thresholds in *NoCommit* ($D = 0.19$, p-value < 0.001).

Table 27 T-Tests: Providers Lying Behavior (*NoCommit*) in Study 3

Sample	Measure	Avg Value	Comparison Benchmark	t(df)	p-value	Bootstrap p-value
All Cases	Lie	0.70	0	$t(6) = 2.70$	0.018	0.015
All Cases	Absolute Lie	1.52	0	$t(6) = 7.40$	< 0.001	< 0.001
Lie $\neq$ 0 Only	Lie	1.42	0	$t(6) = 2.88$	0.014	0.020
Lie $\neq$ 0 Only	Absolute Lie	2.96	0	$t(6) = 9.41$	< 0.001	< 0.001

Notes: The variable *Lie* is defined as $Lie = \theta' - \theta$. Session-level averages are used as the unit of analysis. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 28 OLS Regressions: Service Providers' Lying Behavior in Study 3

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

Notes: 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.

Figure 36 Provider Threshold Distributions by Participant (*Commit* in Study 3)





Figure 37 Provider Threshold Distributions by Participant (*NoCommit* in Study 3)

Implemented (θ) — Communicated (θ') —



E.3.3 Social Welfare We examine aggregate social welfare, both realized and implied. Table 29 presents comparisons between average welfare and relevant theoretical benchmarks. Table 30 compares average welfare between treatments. In the *Commit* treatment, we study the implied welfare conditional on some provider thresholds θ' of interest. Table 31 presents comparisons between average implied welfare in *Commit* conditional on θ' and benchmarks/treatments of interest. In the *NoCommit* treatment, we study the implied social welfare conditional on a provider lying behavior: honest ($\theta' = \theta$), overstate ($\theta' > \theta$), and understate ($\theta' < \theta$). Table 32 presents comparisons between average implied welfare in *NoCommit* conditional on lying behavior and treatments of interest. Table 33 presents OLS regressions regarding social welfare. Table 34 presents comparisons between average implied welfare and the *FullInfo* theoretical benchmark.

Table 29 T-Tests: Social Welfare vs. Theoretical Benchmarks in Study 3

Treatment	Welfare Type	Avg Welfare	Theory	t(df)	p-value	Bootstrap p-value
<i>NoInfo</i>	Realized	172.33	40	$t(2) = 20.20$	0.001	0.034
<i>NoInfo</i>	Implied	172.83	40	$t(2) = 10.18$	0.005	0.074
<i>Commit</i>	Realized	199.37	250	$t(4) = -14.09$	< 0.001	< 0.001
<i>Commit</i>	Implied	200.70	250	$t(4) = -10.15$	< 0.001	< 0.001
<i>NoCommit</i>	Realized	194.61	40	$t(6) = 28.10$	< 0.001	< 0.001
<i>NoCommit</i>	Implied	198.70	40	$t(6) = 41.13$	< 0.001	< 0.001

Notes: Session-level averages are used as the unit of analysis. Theoretical benchmarks are treatment-specific predictions derived from equilibrium analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 30 T-Tests: Treatment Effects, Social Welfare in Study 3

Comparison	Welfare Type	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Welch Bootstrap p-value
<i>NoInfo</i> vs. <i>Commit</i>	Realized	172.33 vs. 199.37	$t(6) = -3.99$	0.003	0.006	0.016	0.019
<i>NoInfo</i> vs. <i>NoCommit</i>	Realized	172.33 vs. 194.61	$t(8) = -2.33$	0.023	0.018	0.024	0.018
<i>Commit</i> vs. <i>NoCommit</i>	Realized	199.37 vs. 194.61	$t(10) = 0.65$	0.263	0.253	0.243	0.249
<i>NoInfo</i> vs. <i>Commit</i>	Implied	172.83 vs. 200.70	$t(6) = -2.42$	0.026	0.028	0.077	0.033
<i>NoInfo</i> vs. <i>NoCommit</i>	Implied	172.83 vs. 198.70	$t(8) = -2.61$	0.015	0.016	0.089	0.036
<i>Commit</i> vs. <i>NoCommit</i>	Implied	200.70 vs. 198.70	$t(10) = 0.32$	0.376	0.371	0.378	0.384
<i>NoInfo</i> vs. <i>Commit</i>	Implied (2nd Half)	181.52 vs. 204.53	$t(6) = -1.78$	0.062	0.062	0.080	0.052
<i>NoInfo</i> vs. <i>NoCommit</i>	Implied (2nd Half)	181.52 vs. 202.55	$t(8) = -2.56$	0.017	0.021	0.093	0.039
<i>Commit</i> vs. <i>NoCommit</i>	Implied (2nd Half)	204.53 vs. 202.55	$t(10) = -0.27$	0.396	0.396	0.409	0.389

Notes: We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates. “2nd Half” refers to observations from the second half of each session.

Table 31 T-Tests: Treatment Effects, Implied Social Welfare by Thresholds θ' of Interest (*Commit*) in Study 3

Comparison	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
<i>Commit</i> ($\theta' = 0$) vs. <i>NoInfo</i>	196.46 vs. 172.83	$t(6) = 1.47$	0.096	0.091	0.108	0.066
<i>Commit</i> ($\theta' = 2$) vs. <i>NoInfo</i>	207.66 vs. 172.83	$t(6) = 2.97$	0.012	0.013	0.050	0.035
<i>Commit</i> ($\theta' = 8$) vs. <i>NoInfo</i>	192.26 vs. 172.83	$t(6) = 1.47$	0.096	0.106	0.138	0.088
<i>Commit</i> ($\theta' = 2$) vs. 250	207.66 vs. 250	$t(4) = -8.27$	< 0.001	0.001	—	—

Notes: We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 32 T-Tests: Treatment Effects, Implied Social Welfare by Lying Behavior (*NoCommit*) in Study 3

Comparison	Avg Welfare	t(df)	p-value	Bootstrap p-value	Welch p-value	Bootstrap Welch p-value
<i>NoCommit(Honest)</i> vs. <i>NoInfo</i>	195.08 vs. 172.83	$t(8) = 2.39$	0.022	0.026	0.113	0.041
<i>NoCommit(Honest)</i> vs. <i>Commit</i>	195.08 vs. 207.66	$t(10) = -1.01$	0.168	0.162	0.182	0.171
<i>NoCommit(Understate)</i> vs. <i>NoInfo</i>	169.06 vs. 172.83	$t(8) = -0.34$	0.368	0.369	0.403	0.383
<i>NoCommit(Understate)</i> vs. <i>Commit</i>	169.06 vs. 207.66	$t(10) = -4.52$	< 0.001	< 0.001	< 0.001	0.001
<i>NoCommit(Overstate)</i> vs. <i>NoInfo</i>	218.91 vs. 172.83	$t(8) = 4.70$	< 0.001	0.002	0.031	0.035
<i>NoCommit(Overstate)</i> vs. <i>Commit</i>	218.91 vs. 207.66	$t(10) = 3.02$	0.006	0.010	0.008	0.016
<i>NoCommit(Overstate)</i> vs. 250	218.91 vs. 250	$t(6) = -8.29$	< 0.001	< 0.001	—	—

Notes: Honest ($\theta' = \theta$), overstate ($\theta' > \theta$), and understate ($\theta' < \theta$). We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

Table 33 OLS Regressions: Social Welfare in Study 3

	Realized Social Welfare					Implied Social Welfare				
	(1a)	(2a)	(3a) [†]	(4a)	(5a) [‡]	(1b)	(2b)	(3b) [†]	(4b)	(5b) [‡]
(Intercept)	146.46*** (9.56)	146.46*** (9.43)	179.23*** (10.65)	146.46*** (9.26)	146.46*** (9.27)	157.80*** (14.57)	157.80*** (14.36)	188.40*** (6.97)	157.80*** (14.10)	157.80*** (14.13)
<i>Commit</i>	35.60* (11.59)	-	2.83 (12.42)	35.60** (11.22)	35.60** (11.24)	33.81* (14.96)	-	3.20 (7.72)	33.81* (14.48)	33.81* (14.51)
<i>NoCommit</i>	-	32.77* (14.30)	-	32.77* (14.05)	-	-	30.61* (15.99)	-	30.61* (15.71)	-
Round	1.26** (0.35)	1.26** (0.35)	0.75* (0.31)	1.26** (0.34)	1.26** (0.34)	0.73** (0.16)	0.73** (0.15)	0.50* (0.20)	0.73*** (0.15)	0.73*** (0.15)
<i>Commit</i> *Round	-0.42 (0.54)	-	0.09 (0.50)	-0.42 (0.52)	-0.42 (0.52)	-0.29 (0.31)	-	-0.06 (0.33)	-0.29 (0.30)	-0.29 (0.30)
<i>NoCommit</i> *Round	-	-0.51 (0.46)	-	-0.51 (0.46)	-	-	-0.23 (0.26)	-	-0.23 (0.25)	-
<i>NoCommit(Honest)</i>	-	-	-	-	48.43** (16.21)	-	-	-	-	33.27** (14.93)
<i>NoCommit(Overstate)</i>	-	-	-	-	60.55** (16.02)	-	-	-	-	52.55** (17.39)
<i>NoCommit(Understate)</i>	-	-	-	-	-27.26* (12.87)	-	-	-	-	2.96 (15.13)
<i>NoCommit(Honest)</i> *Round	-	-	-	-	-1.3* (0.67)	-	-	-	-	-0.57* (0.25)
<i>NoCommit(Overstate)</i> *Round	-	-	-	-	-0.73 (0.52)	-	-	-	-	-0.35 (0.33)
<i>NoCommit(Understate)</i> *Round	-	-	-	-	0.07 (0.58)	-	-	-	-	-0.43 (0.29)
N	640	800	960	1200	1200	640	800	960	1200	1200
R ²	0.04	0.02	0.01	0.02	0.06	0.16	0.11	0.03	0.11	0.21

Notes: 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 *NoInfo* treatment, and take as baseline the *NoCommit* treatment to compare it with the *Commit* treatment.

‡: Models (5a) and (5b) consider the disaggregated *NoCommit* treatment: honest ($\theta' = \theta$), overstate ($\theta' > \theta$), and understate ($\theta' < \theta$).

Table 34 T-Tests: Implied Social Welfare vs. *FullInfo* Theoretical Benchmark in Study 3

Treatment	Avg Welfare	Theory	t(df)	p-value	Bootstrap p-value
<i>NoInfo</i>	172.83	199.6	$t(2) = -2.05$	0.088	0.035
<i>Commit</i>	200.70	199.6	$t(4) = 0.23$	0.416	0.426
<i>Commit</i> ($\theta = 2$)	207.66	199.6	$t(4) = 1.58$	0.095	0.051
<i>NoCommit</i>	198.70	199.6	$t(6) = -0.23$	0.412	0.418
<i>NoCommit(Honest)</i>	195.08	199.6	$t(6) = -1.41$	0.104	0.105
<i>NoCommit(Overstate)</i>	218.91	199.6	$t(6) = 5.15$	0.001	0.009
<i>NoCommit(Understate)</i>	169.06	199.6	$t(6) = -6.41$	< 0.001	0.001

Notes: We use session-level averages as the unit of analysis. See Appendix C for the computation of “Implied” welfare. Alternative hypotheses are set given the direction of the observed effect. Bootstrap p-values computed using 9,999 replicates.

E.3.4 Practice Rounds In the practice rounds, low-need customers arriving at an unobservable join with an average probability of 0.627. Compared with the theoretical prediction of 1: $t(14) = -18.97$, p-value < 0.001, bootstrap t-test p-value < 0.001. Moreover, compared with a flipping a coin behavior with joining probability 0.5: $t(14) = 6.492$, p-value < 0.001, bootstrap t-test p-value < 0.001.

F Post Experimental Survey

In all studies, 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*, *CommitBOT*, *NoCommit*)
- Q5: The service provider's information strategy had an impact on my decisions (*Commit*, *CommitBOT*, *NoCommit*)
- Q6: I trusted the communicated information strategy (*Commit*, *CommitBOT*, *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 how the information strategy works (*Commit*, *NoCommit*)
- Q5: My information strategy had an impact on customers' decisions (*Commit*, *NoCommit*)
- Q6: Customers trusted my communicated information strategy (*Commit*, *NoCommit*)

Tables 35 and 36 present session average results and t-tests for comparisons between treatments in Study 1, 2, and 3, respectively. Table 37 reports the t-test results for cross-study comparisons of post-experimental survey responses, highlighting differences between similar treatment conditions across Studies 1, 2, and 3 for both Customers and Service Providers.

Table 35 Post Experimental Survey Responses in Study 1,2, and T-Tests in Study 1

(a) Average Survey Responses								
Statement	Customers				Service Providers			
	NoInfo	Commit	NoCommit	CommitBot (Study 2)	NoInfo	Commit	NoCommit	
Q1	5.72	5.22	4.91	5.34	5.12	5.38	5.29	
Q2	4.67	4.86	5.04	4.73	2.81	4.69	5.14	
Q3	4.72	5.25	4.66	4.44	5.12	5.31	6.21	
Q4	-	4.97	4.73	5.14	-	5.62	5.57	
Q5	-	5.30	4.82	5.20	-	5.38	6.07	
Q6	-	4.62	4.23	4.97	-	4.81	5.71	
(b) T-Tests (Study 1 only)								
Statement	Customers				Service Providers			
	NoInfo vs. Commit	NoInfo vs. NoCommit	Commit vs. NoCommit		NoInfo vs. Commit	NoInfo vs. NoCommit	Commit vs. NoCommit	
Q1	t(14) = 2.84, p = 0.006	t(13) = 3.46, p = 0.002	t(13) = 1.32, p = 0.105		t(14) = -0.50, p = 0.311	t(13) = -0.32, p = 0.375	t(13) = 0.21, p = 0.418	
Q2	t(14) = -0.92, p = 0.185	t(13) = -1.62, p = 0.064	t(13) = -0.82, p = 0.212		t(14) = -3.31, p = 0.003	t(13) = -4.29, p < 0.001	t(13) = -0.87, p = 0.200	
Q3	t(14) = -2.4, p = 0.015	t(13) = 0.17, p = 0.430	t(13) = 1.76, p = 0.050		t(14) = -0.29, p = 0.387	t(13) = -2.30, p = 0.019	t(13) = -1.44, p = 0.086	
Q4	-	-	t(13) = 0.71, p = 0.243		-	-	t(13) = 0.08, p = 0.467	
Q5	-	-	t(13) = 1.36, p = 0.097		-	-	t(13) = -1.22, p = 0.120	
Q6	-	-	t(13) = 1.35, p = 0.099		-	-	t(13) = -1.37, p = 0.096	

Table 36 Post Experimental Survey Responses and T-Tests in Study 3

(a) Average Survey Responses						
Statement	Customers			Service Providers		
	NoInfo	Commit	NoCommit	NoInfo	Commit	NoCommit
Q1	4.69	4.45	4.54	4.83	4.20	4.07
Q2	3.10	3.38	3.22	2.33	3.60	3.21
Q3	3.21	3.55	3.43	3.33	3.10	3.86
Q4	-	4.39	3.94	-	4.10	5.00
Q5	-	5.65	4.41	-	4.60	4.64
Q6	-	3.98	3.28	-	2.60	4.14
(b) T-Tests						
Statement	Customers			Service Providers		
	NoInfo vs. Commit	NoInfo vs. NoCommit	Commit vs. NoCommit	NoInfo vs. Commit	NoInfo vs. NoCommit	Commit vs. NoCommit
Q1	t(6) = 2.37, p = 0.027	t(8) = 0.45, p = 0.329	t(10) = -0.38, p = 0.353	t(6) = 1.35, p = 0.112	t(8) = 1.23, p = 0.125	t(10) = 0.27, p = 0.396
Q2	t(6) = -0.62, p = 0.278	t(8) = -0.40, p = 0.347	t(10) = 0.47, p = 0.322	t(6) = -1.10, p = 0.156	t(8) = -1.23, p = 0.126	t(10) = 0.68, p = 0.253
Q3	t(6) = -0.84, p = 0.216	t(8) = -1.18, p = 0.136	t(10) = 0.45, p = 0.330	t(6) = 0.35, p = 0.367	t(8) = -0.81, p = 0.221	t(10) = -1.59, p = 0.071
Q4	-	-	t(10) = 2.46, p = 0.017	-	-	t(10) = -1.95, p = 0.040
Q5	-	-	t(10) = 3.25, p = 0.004	-	-	t(10) = -0.079, p = 0.469
Q6	-	-	t(10) = 1.91, p = 0.042	-	-	t(10) = -4.13, p = 0.001

Table 37 T-Tests: Cross-Study Comparisons of Post Experimental Survey Responses

Statement	Customers					Service Providers		
	NoInfo (S1 vs. S3)	Commit (S1 vs. S2)	Commit (S1 vs. S3)	Commit (S2 vs. S3)	NoCommit (S1 vs. S3)	NoInfo (S1 vs. S3)	Commit (S1 vs. S3)	NoCommit (S1 vs. S3)
Q1	t(9) = 4.88, p < 0.001	t(14) = -0.55, p = 0.291	t(11) = 4.53, p < 0.001	t(11) = 3.53, p = 0.002	t(12) = 1.28, p = 0.111	t(9) = 0.41, p = 0.343	t(11) = 2.64, p = 0.011	t(12) = 2.68, p = 0.010
Q2	t(9) = 5.55, p < 0.001	t(14) = 0.52, p = 0.307	t(11) = 5.09, p < 0.001	t(11) = 3.92, p = 0.001	t(12) = 7.73, p < 0.001	t(9) = 0.52, p = 0.307	t(11) = 1.57, p = 0.072	t(12) = 4.98, p < 0.001
Q3	t(9) = 5.64, p < 0.001	t(14) = 2.19, p = 0.022	t(11) = 5.53, p < 0.001	t(11) = 1.84, p = 0.046	t(12) = 3.84, p < 0.001	t(9) = 2.51, p = 0.017	t(11) = 3.06, p = 0.005	t(12) = 5.47, p < 0.001
Q4	-	t(14) = -0.55, p = 0.296	t(11) = 1.77, p = 0.052	t(11) = 2.91, p = 0.007	t(12) = 3.17, p = 0.004	-	t(11) = 2.25, p = 0.023	t(12) = 1.13, p = 0.140
Q5	-	t(14) = 0.29, p = 0.386	t(11) = -1.25, p = 0.119	t(11) = -1.34, p = 0.103	t(12) = 0.97, p = 0.173	-	t(11) = 1.12, p = 0.142	t(12) = 3.26, p = 0.003
Q6	-	t(14) = -1.22, p = 0.120	t(11) = 1.92, p = 0.041	t(11) = 2.95, p = 0.006	t(12) = 3.07, p = 0.005	-	t(11) = 3.36, p = 0.003	t(12) = 3.17, p = 0.003

In here, we present a list of findings from the post-experimental survey:

- *Evidence of increased conflict of interest in Study 3.* Post-experimental survey results support the notion that study 3 presents a more salient conflict of interest. Specifically, responses to item Q3, which measures perceived conflict of interest, show a clear and statistically significant increase across studies: participants in Study 1 reported the least conflict, followed by Study 2, with Study 3 showing the highest perceived conflict. This upward trend held consistently across customers, providers, and all treatment conditions (see Table 37). A similar pattern emerged for item Q6, which measures perceived credibility in the communicated signaling policy. Credibility ratings were highest in Study 1 and 2, and significantly lower in Study 3—for both customers evaluating providers' credibility and providers evaluating their own credibility—again consistently across treatments (see Table 37).
- *Evidence of increased complexity in Study 3.* Post-experimental survey results support the notion that study 3 presents a more complex decision environment. Specifically, responses to item Q2, which measures perceived control over outcomes, show a clear and statistically significant pattern across studies. Participants' perceived control was highest in Study 1, remained largely unchanged in Study 2, and was lowest in Study 3. This pattern held consistently across all treatments for customers, and was also evident for the *Commit* and *NoCommit* treatments among service providers (see Table 37). Similarly, responses to item Q4, which measures perceived comprehension difficulty (regarding signaling), show a statistically significant pattern across studies. Among customers, perceived difficulty was lowest in Study 1, remained largely unchanged in Study 2, and was highest in Study 3. This pattern held consistently across all customer treatment conditions. Among service providers, a similar increasing pattern in perceived difficulty was observed across studies, though it was only marginally significant ($p \approx 0.15$) in the *NoCommit* condition (see Table 37).
- *Credibility effects.* Post-experimental survey results support the idea that perceived credibility differs depending on commitment power. Specifically, item Q6—which measures perceived credibility in the communicated signaling—shows that for customers, the provider's credibility is statistically higher under the *Commit* condition than under *NoCommit* in both Study 1 (at the 10% significance level) and Study 3. Interestingly, for providers, their own perceived credibility is statistically higher under the *NoCommit* condition than under *Commit* in both Study 1 (at the 10% significance level) and Study 3 (see Tables 35 and 36).
- *Influence effects.* Post-experimental survey results support the idea that perceived decision influence differs depending on commitment power. Specifically, item Q5—which measures perceived influence on decisions—indicates that customers more strongly believed their decisions were influenced under the *Commit* condition than under *NoCommit* in both Study 1 (at the 10% significance level) and Study 3 (see Tables 35 and 36).

G Behavioral Structural Model: 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):

$$\mathbb{P}_i(a_i) = \frac{e^{\mathbb{E}[\pi_i(a_i, \mathbf{A}_{-i})]/\beta_i}}{\sum_{a'_i \in S_i} e^{\mathbb{E}[\pi_i(a'_i, \mathbf{A}_{-i})]/\beta_i}}. \quad (3)$$

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 (3), 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 $\beta_i \rightarrow 0$, player i chooses the payoff-maximizing alternative with certainty (i.e., the theoretical prediction under full rationality). At the other extreme, when $\beta_i \rightarrow \infty$, 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. For simplicity, we assume that high-need customers always join the system, and thus focus entirely on low-need customer behavior. This is consistent with our experimental results in Study 1 and 2, and with the experimental design in Study 3.

G.1 Estimation Procedure

G.1.1 NoInfo Treatment In this case, low-need customers join with probability:

$$\varphi(\beta) = \frac{e^{(r_L - c(\mathbb{E}[Q] + 1))/\beta}}{e^{(r_L - c(\mathbb{E}[Q] + 1))/\beta} + e^{v_L/\beta}}. \quad (4)$$

In this expression, customers compute the expected number 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 behavior. That is, the QRE of the game represents the solution of the fixed point problem given by equation (4), 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 $t \in \{1, \dots, T\}$ and for each cohort $c \in \{1, \dots, C\}$, such data set contains for each customer $i \in \{1, \dots, \Lambda\}$, 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-likelihood function is given by:

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

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.

G.1.2 FullInfo Treatment In this case, low-need customers join with probability:

$$\varphi(q; \beta) = \frac{e^{(r_L - c(q+1))/\beta}}{e^{(r_L - c(q+1))/\beta} + e^{v_L/\beta}}, \quad \text{for } q = 0, \dots, \Lambda - 1 \quad (5)$$

In this expression, customers observe the number of customers in the system q upon their arrival. Since customer decisions are based on the observed queue length, $\varphi(q; \beta)$ directly characterizes the QRE of the game, denoted by $\varphi^*(q; \beta)$.

Estimation of Relevant Parameter. 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 for each customer $i \in \{1, \dots, \Lambda\}$, their strategy $A_{itc} \equiv \{a_{itc}(q) \in \{0, 1\} \forall q \in \{0, \dots, \Lambda - 1\}\}$, where 1 represents the *Join* action and 0 the *Balk* action. Based on this, the log-likelihood function is given by:

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

G.1.3 Commit Treatment In this case, given an implemented threshold θ and a signal ς , low-need customers join with probability:

$$\varphi_s(\theta; \beta) = \frac{e^{(r_L - c(\mathbb{E}[Q|\varsigma=s, \theta]+1))/\beta}}{e^{(r_L - c(\mathbb{E}[Q|\varsigma=s, \theta]+1))/\beta} + e^{v_L/\beta}} \quad \text{for } \theta = 1, \dots, \Lambda, \quad (6)$$

$$\varphi_l(\theta; \beta) = \frac{e^{(r_L - c(\mathbb{E}[Q|\varsigma=l, \theta]+1))/\beta}}{e^{(r_L - c(\mathbb{E}[Q|\varsigma=l, \theta]+1))/\beta} + e^{v_L/\beta}} \quad \text{for } \theta = 0, \dots, \Lambda - 1. \quad (7)$$

In these expressions, customers compute the expected number 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 :

$$\vartheta(\theta; \beta, \beta_m) = \frac{e^{(\Omega(\varphi_s(\theta; \beta), \varphi_l(\theta; \beta), \theta))/\beta_m}}{\sum_{\theta} e^{(\Omega(\varphi_s(\theta; \beta), \varphi_l(\theta; \beta), \theta))/\beta_m}} \quad \text{for } \theta = 0, \dots, \Lambda. \quad (8)$$

Based on the above, in equilibrium, customers' and the service provider's expectations and beliefs are consistent with customers' and the service provider's behavior. That is, the QRE of the game is given by the solution of the system of equations (6) - (8), denoted by $\varphi_s^*(\theta; \beta)$ for $\theta \in \{1, \dots, \Lambda\}$, $\varphi_l^*(\theta; \beta)$ for $\theta \in \{0, \dots, \Lambda - 1\}$ and $\vartheta^*(\theta; \beta, \beta_m)$ for $\theta \in \{0, \dots, \Lambda\}$. 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 (6) and (7). Then, based on such customer equilibrium, we can compute $\vartheta^*(\theta; \beta, \beta_m)$ by plugging the customer equilibrium directly in (8). 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:

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

where q represents the current number of customers in queue, and w the remaining customers to arrive. 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:

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

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 signaling 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 Ω . 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, $p_H(r_H - c(q+1)) + (1-p_H)(r_L - c(q+1))(\sigma_\theta(q, w)\varphi_s(\theta) + (1-\sigma_\theta(q, w))\varphi_l(\theta)) + v_L(1-p_H)(\sigma_\theta(q, w)(1-\varphi_s(\theta)) + (1-\sigma_\theta(q, w))(1-\varphi_l(\theta)))$, 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:

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

with boundary conditions $V(q, 0|\theta) = p_H(r_H - c(q+1)) + (1-p_H)(r_L - c(q+1))(\sigma_\theta(q, 0)\varphi_s(\theta) + (1-\sigma_\theta(q, 0))\varphi_l(\theta)) + v_L(1-p_H)(\sigma_\theta(q, 0)(1-\varphi_s(\theta)) + (1-\sigma_\theta(q, 0))(1-\varphi_l(\theta)))$ 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 signaling 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 $\theta \in \{1, \dots, \Lambda\}$, $\varphi_l^*(\theta; \beta)$ for $\theta \in \{0, \dots, \Lambda-1\}$ and $\vartheta^*(\theta; \beta, \beta_m)$ for $\theta \in \{0, \dots, \Lambda\}$. 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} . Moreover, the data set contains for each customer $i \in \{1, \dots, \Lambda\}$, their strategy $A_{itc} \equiv \{a_{itc}(s = s|\theta_{jtc}), a_{itc}(s = l|\theta_{jtc})\}$, where $a_{itc} \in \{0, 1\}$ is such that 1 represents the *Join* action and 0 the *Balk* action. Based on this, the log-likelihood function is given by:

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

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.

G.1.4 NoCommit Treatment In this case, given a communicated threshold θ' and a signal ς , low-need customers join with probability:

$$\varphi_s(\theta'; \beta, \beta_m, \kappa) = \frac{e^{(r_L - c(\mathbb{E}[Q|\varsigma=s, \theta'] + 1))/\beta}}{e^{(r_L - c(\mathbb{E}[Q|\varsigma=s, \theta'] + 1))/\beta} + e^{v_L/\beta}} \quad \text{for } \theta' = 0, \dots, \Lambda, \quad (9)$$

$$\varphi_l(\theta'; \beta, \beta_m, \kappa) = \frac{e^{(r_L - c(\mathbb{E}[Q|\varsigma=l, \theta'] + 1))/\beta}}{e^{(r_L - c(\mathbb{E}[Q|\varsigma=l, \theta'] + 1))/\beta} + e^{v_L/\beta}} \quad \text{for } \theta' = 0, \dots, \Lambda. \quad (10)$$

In these expressions, customers compute the expected number 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 (see below). 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. Because providers have the possibility to lie (θ and θ' need to be the same), consistent with the notion of preferences for truth-telling in the behavioral economics literature (Rosenbaum et al. 2014, Abeler et al. 2019), we consider the possibility that providers incur 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). 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):

$$\vartheta(\theta, \theta'; \beta, \beta_m, \kappa) = \frac{e^{(\Omega(\varphi_s(\theta'), \varphi_l(\theta'), \theta) - \kappa|\theta' - \theta|)/\beta_m}}{\sum_{\theta} \sum_{\theta'} e^{(\Omega(\varphi_s(\theta'), \varphi_l(\theta'), \theta) - \kappa|\theta' - \theta|)/\beta_m}} \quad \text{for } \theta = 0, \dots, \Lambda \quad \text{and} \quad \theta' = 0, \dots, \Lambda. \quad (11)$$

Based on the above, in equilibrium, customers' and the service provider's expectations and beliefs are consistent with customers' and the service provider's behavior. That is, the QRE of the game is given by the solution of the system of equations (9) - (11), denoted by $\varphi_s^*(\theta'; \beta, \beta_m, \kappa)$ for $\theta' \in \{0, \dots, \Lambda\}$, $\varphi_l^*(\theta'; \beta, \beta_m, \kappa)$ for $\theta' \in \{0, \dots, \Lambda\}$ and $\vartheta^*(\theta, \theta'; \beta, \beta_m, \kappa)$ for $\theta \in \{0, \dots, \Lambda\}$ and $\theta' \in \{0, \dots, \Lambda\}$. 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:

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

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 signaling 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 signaling 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 $\mathbb{P}(\theta | \theta') = \frac{\vartheta(\theta, \theta')}{\vartheta(\theta')}$, 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, $p_H(r_H - c(q + 1)) + (1 - p_H)(r_L - c(q + 1))(\sigma_\theta(q, w)\varphi_s(\theta') + (1 - \sigma_\theta(q, w))\varphi_l(\theta')) + v_L(1 - p_H)(\sigma_\theta(q, w)(1 - \varphi_s(\theta')) + (1 - \sigma_\theta(q, w))(1 - \varphi_l(\theta')))$, 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:

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

with boundary conditions $V(q, 0 | \theta, \theta') = p_H(r_H - c(q + 1)) + (1 - p_H)(r_L - c(q + 1))(\sigma_\theta(q, 0)\varphi_s(\theta') + (1 - \sigma_\theta(q, 0))\varphi_l(\theta')) + v_L(1 - p_H)(\sigma_\theta(q, 0)(1 - \varphi_s(\theta')) + (1 - \sigma_\theta(q, 0))(1 - \varphi_l(\theta')))$ 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 signaling 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 $\theta' \in \{0, \dots, \Lambda\}$, $\varphi_l^*(\theta'; \beta, \beta_m, \kappa)$ for $\theta' \in \{0, \dots, \Lambda\}$ and $\vartheta^*(\theta, \theta'; \beta, \beta_m, \kappa)$ for $\theta \in \{0, \dots, \Lambda\}$ and $\theta' \in \{0, \dots, \Lambda\}$. 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 $i \in \{1, \dots, \Lambda\}$, 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-likelihood function is given by:

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

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.

G.2 Results

For each treatment, we estimate the set of relevant parameters, β, β_m, κ . The estimated QRE models are presented in Tables 38 and 39. These tables include 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 provide the best fit for our data. Figures 38 and 39 present customer and provider’s choice predictions under these full models. We discuss the results from these models next.

Table 38 QRE Estimation Results (standard errors in parentheses) in Study 1 and 2

	<i>NoInfo</i>	<i>FullInfo (Study 1)</i>	<i>FullInfo (Study 2)</i>	<i>Commit</i>		<i>CommitBOT (Study 2)</i>	<i>NoCommit</i>			
	(1) Full	(2) Full	(3) Full	(4a) $\beta = \beta_m$	(4b) Full	(5) Full	(6a) $\beta = \beta_m, \kappa = 0$	(6b) $\beta = \beta_m$	(6c) $\kappa = 0$	(6d) Full
$\hat{\beta}$	226.59* (183.40)	21*** (0.45)	17.20*** (0.64)	16.74*** (0.04)	15.76*** (0.06)	14.68*** (0.07)	26.71*** (0.53)	18.22*** (0.34)	17.27*** (0.54)	15.63*** (0.44)
$\hat{\beta}_m$	-	-	-	-	37.93*** (10.73)	-	-	-	115.47*** (20.72)	60.88*** (8.92)
$\hat{\kappa}$	-	-	-	-	-	-	-	13.72*** (0.05)	-	41.10*** (6.23)
Obs. [Customers/Providers]	2560/-	1840/-	640/-	2560/640		-	2240/560			
Log-Likelihood	-1773.70	-2460.01	-708.09	-2304.30	-2295.27	-1561.55	-3849.9	-3167.77	-3723.45	-3078.12
LL-ratio test (vs. Full)	-	-	-	$\chi^2(1) = 18.06***$	-	-	$\chi^2(2) = 1543.57***$	$\chi^2(1) = 179.32***$	$\chi^2(1) = 1290.66***$	-
Expected Welfare Ω	617.48	619.04	619.28	615.62	614.75	621.85	613.16	616.07	609.27	613.24

Table 39 QRE Estimation Results (standard errors in parentheses) in Study 3

	<i>NoInfo</i>	<i>Commit</i>		<i>NoCommit</i>			
	(1) Full	(2a) $\beta = \beta_m$	(2b) Full	(3a) $\beta = \beta_m, \kappa = 0$	(3b) $\beta = \beta_m$	(3c) $\kappa = 0$	(3d) Full
$\hat{\beta}$	91.25*** (16.65)	36.80*** (0.07)	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/-	3,200/400		4,480/560			
Log-Likelihood	-1,311.51	-3,417.96	-3,380.37	-7,294.67	-7,064.32	-7,246.51	-7,052.61
LL-ratio test (vs. Full)	-	$\chi^2(1) = 75.17***$	-	$\chi^2(2) = 484.11***$	$\chi^2(1) = 23.41***$	$\chi^2(1) = 387.78***$	-
Expected Welfare Ω	176.95	174.99	179.44	175.04	176.24	168.09	174.19

The estimated QRE models yield a strong quantitative and qualitative fit to the experimental data, capturing key patterns of behavior across treatments and studies. In both Study 1 and Study 3 (see Tables 38 and 39), we find that the estimated bounded rationality parameters $\hat{\beta}$ and $\hat{\beta}_m$ are consistently higher under the *NoCommit* treatment relative to *Commit*, suggesting greater decision noise in the *NoCommit* environment. This is consistent with the interpretation that *NoCommit* introduces additional strategic complexity, due to the possibility of non-binding and potentially deceptive communication, which QRE captures via increased stochasticity in players’ responses.

Interestingly, the highest $\hat{\beta}$ values are observed in the *NoInfo* condition, where customers receive no signals. In this setting, joining decisions appear close to uniform. This suggests that the absence of information represents a highly complex decision environment for customers. This points to the fact that any type of information - even non-binding information - can significantly reduce complexity and guide behavior, as evidenced by the lower $\hat{\beta}$ values in the signaling treatments. At the opposite end, the lowest $\hat{\beta}$ is found in *CommitBOT*, where provider thresholds are automatically generated (see Table 38). This is intuitive since such treatment effectively removes all the provider’s noise from the customers’ decision process.

Beyond parameter estimates, the QRE model captures dynamic behavioral responses to changes in the strategic environment. In both *Commit* and *NoCommit*, we observe that the probability of joining decreases with the communicated threshold θ' under short signals and long signals, consistent with shifts in expected payoffs (Figures 38a, 38c, 39a, 39c). The QRE model replicates these trends, predicting directionally correct comparative statics in joining behavior as a function of θ' , conditional on signal.


Furthermore, the attenuated sensitivity of joining behavior to θ' in the *NoCommit* condition is well-explained by QRE (see Figures 38c and 39c). Under this treatment, due to the possibility of misrepresentation by providers, customers cannot fully infer the true threshold θ from θ' . However, there is partial informativeness because a positive lying cost estimated parameter $\hat{\kappa} > 0$, penalizes message-deviation and endogenously limits strategic deception. This partial informativeness flattens the expected payoff gradients, reducing the responsiveness of the quantal response. We find that introducing a lying cost in the QRE model is sufficient to account for this behavior, generating both non-uniform reporting among providers and meaningful inference by customers.

The QRE predictions for provider threshold choices successfully recovers a unimodal distribution of selected thresholds across treatments, with predicted modes closely matching empirical modes (Figures 38b, 38d, 39b, 39d). Overall, these results indicate that QRE provides a coherent, behaviorally grounded framework that jointly explains both customer and provider behavior across varying informational and strategic contexts.


Finally, we note that the QRE framework captures the observed deviations in social welfare from the theoretical predictions. These deviations are shaped by bounded rationality: as the level of bounded rationality increases, the welfare differences across treatments shrink (see Figure 40). Specifically, in the case of the *NoInfo* treatment, such deviations are welfare-enhancing, while for the rest treatments bounded rationality reduces welfare. As a result, the theoretically predicted welfare gains from information provision are substantially dampened under increased bounded rationality. Indeed, as shown in Tables 38 and 39, the welfare improvements from communication-based treatments are minimal, if present at all.

That said, we offer a note of caution: the QRE-based welfare predictions are slightly *biased* in *Commit*, *CommitBOT* and *NoCommit* treatments, as the model tends to *over-predict* the joining rates of customers (in such treatments) relative to what is observed in the data. Additionally, the predicted distributions of provider choices are smoother than the empirical ones (see Figures 38 and 39). As a consequence, the estimated QRE presents an implied social welfare in *Commit*, *CommitBOT* and *NoCommit* treatments that is lower than what is observed empirically. This also means that the welfare improvements from communication-based treatments under QRE in Tables 38 and 39 are conservative.

Figure 38 Observations vs. QRE Predictions in Study 1,2 (Full Models from Table 38)




Note: In *NoInfo*, the observed joining probability is $\mathbb{P}(\text{Join}) = 0.512$; under estimated QRE, the prediction is the same. In *CommitBOT (Study 2)*, the observed joining probabilities are $\mathbb{P}(\text{Join} | \text{short}) = 0.977$ and $\mathbb{P}(\text{Join} | \text{long}) = 0.171$; predicted probabilities under QRE are $\mathbb{P}(\text{Join} | \text{short}) = 0.993$ and $\mathbb{P}(\text{Join} | \text{long}) = 0.276$.


Figure 39 Observations vs. QRE Predictions in Study 3 (Full Models from Table 39)

Note: In *NoInfo*, the observed joining probability is $\mathbb{P}(\text{Join}) = 0.571$; under estimated QRE, the prediction is the same.


Figure 40 Comparative Statistics: Implied Social Welfare and QRE Parameters




(a) Study 1: Provider $\beta_m \rightarrow 0$ (*Commit*);
 $\beta_m \rightarrow 0, \kappa = 0$ (*NoCommit*)



(b) Study 1: Provider $\beta_m = 37.93$ (*Commit*);
 $\beta_m = 60.88, \kappa = 41.10$ (*NoCommit*)



(c) Study 3: Provider $\beta_m \rightarrow 0$ (*Commit*);
 $\beta_m \rightarrow 0, \kappa = 0$ (*NoCommit*)



(d) Study 3: Provider $\beta_m = 13.67$ (*Commit*);
 $\beta_m = 113.18, \kappa = 44.21$ (*NoCommit*)