Truth-Telling in a Priority Pricing Mechanism

Prakriti Thami*

November 4, 2024

Abstract

This paper studies the impact of truth-telling preferences on aggregate consumer welfare within a priority pricing (PP) mechanism. Traditional models assume individuals always misrepresent private information to maximize payoffs, yet recent evidence suggests there may be an innate preference for truth-telling. By incorporating these preferences, I find that PP enhances welfare over uniform pricing only when the probability of non-truthful individuals surpasses a critical threshold, suggesting that PP may benefit populations with low truth-telling tendencies but reduce welfare when this tendency is high. To empirically test this, I conducted an online experiment, finding that while PP incentivized truth-telling, its impact did not vary significantly across groups with differing truth-telling tendencies. Instead, participants' beliefs about others' truthfulness emerged as key in shaping behavior. These findings underscore that PP's welfare-enhancing potential depends not only on incentives created by the pricing structure but also on the population's truth-telling tendencies and beliefs, offering valuable insights for designing effective pricing mechanisms.

Keywords: priority pricing, consumer welfare, truth-telling behavior.

JEL Classification: D82; D9, D47, D61

1 Introduction

In many situations where the demand for a service exceeds supply, prioritizing service recipients based on their need or valuation can significantly enhance overall welfare, particularly when timely access to a service is critical. However, when need or valuation is private information, efficiently allocating services becomes more challenging. In such situations, priority pricing (PP), a type of third degree price discrimination, often emerges as a practical method for aligning incentives and ensuring that those who need or value the service the most receive

*I am deeply grateful to my supervisors, Tommy Andersson and Pol Campos-Mercade, for their guidance and invaluable feedback throughout this project. I also wish to thank Erik Wengström, Roel van Veldhuizen, Eva Ranehill, and Adrian Nieto Castro for their insightful comments, which significantly enriched this work. I gratefully acknowledge financial support from the Jan Wallanders och Tom Hedelius Stiftelse and the Stiftelsen För Främjande av Ekonomisk Forskning vid Lund Universitet.

it first. This approach utilizes prices to get service recipients to self-select into distinct priority classes, with those in higher-priority groups experiencing shorter wait times in exchange for higher fees. Examples of this can be seen in real-world contexts such as expedited passport processing,[1] fast-track visa services,[2] and priority medical care,[3] where those who value the service most, or need it most urgently, can gain quicker access by paying a premium.

An implicit assumption underlying this approach to allocation in the face of informational asymmetry is that individuals will invariably misrepresent private information if doing so maximizes their personal gains. However, a large body of experimental evidence indicates otherwise (see, for example, Abeler et al., 2014; Fischbacher and Föllmi-Heusi, 2013; Gneezy, 2005; Gneezy et al., 2018; Hurkens and Kartik, 2009; Lundquist et al., 2009). Further, a meta-analysis of experimental studies on truth-telling provides strong empirical evidence that truth-telling behavior is motivated not only by a desire to be perceived as honest but also by an intrinsic preference for truth-telling (Abeler et al., 2019).

This evidence indicates that if consumer welfare enhancement is the primary objective, then any evaluation of the benefits of adopting a PP scheme must also take into account the truth-telling behavior among individuals. The conventional rationale for and benefits from PP—its ability to align incentives and prevent misrepresentation—becomes less compelling if a substantial proportion of the population is already intrinsically motivated to tell the truth.

While there is a vast body of literature focused on determining the optimal PP scheme under different settings (see, for example, Adiri and Yechiali, 1974; Afeche and Mendelson, 2004; Dolan, 1978; Ghanem, 1975; Mendelson and Whang, 1990; Naor, 1969) and a few studies that examine their welfare effects (Chao & Wilson, 1987; Gershkov & Winter, 2023; Wilson, 1989), the consideration of truth-telling behavior and empirical evidence appears to be largely absent from the existing research.

Accordingly, this paper contributes to the theoretical and practical understanding of PP schemes by integrating preferences for truth-telling into the theoretical framework and providing empirical evidence on the behavioral responses to PP. Specifically, the paper first

[1] For example, in the United States, expedited passport processing is available for an additional \$60, reducing the wait time from 6-8 weeks to 2-3 weeks (U.S. Department of State, n.d.). Similarly, the United Kingdom offers a 1-day Premium service for £207.50 and a 1-week Fast Track service for £178.50, compared to standard application fees ranging from £88.50 to £112.(GOV.UK, n.d.). Several EU countries, including Belgium, Germany, and the Netherlands, also provide expedited passport services for an additional fee (Euronews, 2024).

[2] Many countries, including the United States and United Kingdom, offer expedited processing for select types of visa applications. For example, the UK provides a Priority Service for £500, delivering a decision within 5 days, and a Super Priority Service for £1000, with a decision by the next working day (GOV.UK, 2024). In the US, premium processing is available for certain visa categories, offering expedited service with added fees ranging from \$1,685 to \$2,805 (U.S. Citizenship and Immigration Services, n.d.).

[3] Examples of priority medical care include expedited specialist appointments, fast-track options in emergency departments, priority scheduling for elective surgeries, and concierge medicine, which provides same-day appointments and 24/7 access to doctors for a higher fee. These services aim to reduce wait times for those who can afford them, but they also raise ethical concerns about equity and access. Further, in many countries with mixed healthcare systems, including for example, Australia, Germany, Denmark and Poland, patients can opt to pay for private healthcare services or private insurance and gain quicker access to medical care than what is typically available through public healthcare systems (Australian Government Department of Health and Aged Care, n.d.; Kuchinke et al., 2009; OECD, 2020).

evaluates the expected aggregate consumer welfare of an incentive compatible PP scheme, relative to both free-of-charge and uniform pricing schemes, while accounting for the presence of always truth-telling individuals within a simplified model. The model examines the allocation of two appointment slots between two agents, each with a privately known level of need—either high or low. The analysis identifies unique critical thresholds in the proportion of low-need, non-truthful individuals, beyond which PP enhances expected aggregate consumer welfare relative to the other pricing schemes.

These findings have two key implications: First, PP is unlikely to enhance welfare unless there is a sufficiently high propensity for non-truthfulness within the population. Second, while PP may enhance welfare in populations with a high underlying propensity for non-truthfulness, it could be detrimental in populations where this propensity is low. Further, since the primary mechanism through which the PP scheme impacts expected aggregate consumer welfare is by incentivizing individuals to truthfully report their level of need, it is crucial to empirically test whether the pricing incentive has the desired effect on truth-telling behavior.

To test these theoretical insights, the paper conducts a pre-registered online survey-based experiment designed to assess the behavioral response to an incentive compatible PP scheme. The experiment addresses the challenge of creating groups with differing underlying propensities for truthfulness by first observing participants' tendencies toward non-truthfulness and then using these observations to exogenously shift participant beliefs about the truthfulness propensity within their assigned group. Additionally, these observations are used to weight the data, enabling the simulation of groups with high and low underlying truth-telling propensities for the analysis. Participants were randomly assigned to one of four groups: two treatment groups with priority pricing (truthful and non-truthful) and two control groups without (truthful and non-truthful). The survey was structured into three sections, requiring participants to make a total of six incentivized decisions, two per section.

In Section 1, participants' propensity for non-truthfulness was assessed by assigning them randomly as either Type A or Type B individuals. They were offered a \$1.50 payment for reporting as Type A and \$0 for reporting as Type B, regardless of their actual assigned type. Participants then answered the question "What is your type?" for each possible assignment.

In Section 2, participants simulated a queuing scenario to book a doctor's appointment, with two available slots: one immediate and one for the next day. Participants were assigned as either Type U (Urgent), with higher valuations for the slots (\$10 and \$5), or Type N (Non-Urgent), with lower valuations (\$3 and \$1.50). In the treatment groups, a \$0.75 priority fee was introduced for those who reported as Type U. Participants were informed that the first appointment slot would be given to the participant reporting as Type U, with random allocation if both reported the same type. Participants made decisions once for each possible type assignment.

In Section 3, participants made decisions for the simulated queuing scenario again but were informed that they would be randomly assigned to a group with 9 other participants and could qualify for extra payment based on the characteristics of their assigned group. Specifically, those in the Truthful groups could earn extra payment from their response in the section if 9 out of 10 group members had been truthful in Section 1, while those in the Non-Truthful groups could earn extra if 9 out of 10 members had been non-truthful in Section 1.

The results indicate that while PP reduces the proportion of low-need participants who misrepresent their types, the pricing incentives may not completely eliminate misrepresentation, as 26.7% of participants in the priority pricing treatment groups still reported as being Type U when assigned as Type N. This suggests that the threshold proportion of low-need, non-truthful individuals required for PP to improve welfare could be significantly higher than even the theoretical prediction.

Additionally, contrary to the initial hypothesis, the experiment does not find significant evidence of a differential treatment effect of PP between the truthful and non-truthful groups. This absence of a differential effect appear to be driven by participants adjusting their behavior based on their beliefs about the truthfulness of their group members. Specifically, when participants were informed that they would qualify for extra payment from Section 3 only if 9 out of 10 group members were non-truthful in Section 1, they were more likely to respond non-truthfully themselves in Section 3. Conversely, when prompted with the truthful condition, they tended to respond more truthfully. In other words, participants' beliefs about the truthfulness of others significantly influenced the effectiveness of the pricing incentive. These findings underscore the importance for policymakers and economists to consider both individual and collective behaviors when optimizing PP strategies in practice.

The paper is organized as follows. Section 2 presents a review of the relevant literature. Section 3 presents the theoretical framework and hypothesis that guide the experimental analysis. Section 4 details the experimental design. Section 5 presents the results. Finally, Section 6 discusses the broader implications of the results and offers some concluding insights.

2 Related Literature

The theoretical framework of this paper relates to the literature on incentive-compatible priority pricing and its properties. Naor (1969) formally introduced the use of pricing to manage queues and enhance social welfare, while Dolan (1978) initiated the mechanism design literature on queuing by proposing a scheme where recipients are charged a price equal to the marginal delay cost they impose on others, thereby incentivizing them to truthfully report their actual delay costs. Since these seminal works, much of the focus in this strand of literature has been on determining the optimal priority pricing schemes across various setting. Chapter 4 in Hassin and Haviv (2003) provides a comprehensive overview of the key theoretical developments.

Several papers have also explored the welfare implications of priority pricing schemes. For instance, in markets subject to random shocks, such as electricity markets, Chao and Wilson (1987) shows that priority pricing can achieve the same welfare-maximizing allocation as spot pricing but with lower costs and greater efficiency. Wilson (1989) builds on this by proving that there exists a priority pricing scheme that maximizes total welfare and can be adjusted to redistribute revenues to consumers, leading to a Pareto improvement over random assignment. However, the modeling approach in these papers differ from the current study. They employ dynamic continuous-time models and compare to spot pricing, while this paper uses a simplified static model for tractability and clarity, focusing on comparisons with free-of-charge and uniform pricing. Additionally, while those analyses pertain to total welfare, the present study specifically focuses on consumer welfare.

While Gershkov and Winter (2023) examine the impact of priority pricing on consumer surplus, their focus is on monopoly settings. They find that although priority pricing can improve efficiency, it may reduce consumer surplus due to monopolistic surplus extraction. However, these negative effects can be mitigated if priority services attract new consumers and expand market coverage. A key distinction between these studies and the current paper is that here, prices are used exclusively as incentives for truthful information revelation, with the primary objective being the maximization of expected aggregate consumer welfare.

This paper also relates to the literature on preferences for truth-telling (also commonly referred to as lying aversion), a well-documented phenomenon in experimental economics (see, for example, Abeler et al., 2014; Erat and Gneezy, 2012; Fischbacher and Föllmi-Heusi, 2013; Gneezy, 2005; Gneezy et al., 2018; Hurkens and Kartik, 2009; Lundquist et al., 2009). Studies have shown that truth-telling, even at the expense of self-interest, can be driven by social preferences (Gneezy, 2005) and reputational concerns (Abeler et al., 2019; Gneezy et al., 2018), yet individuals also possess an inherent preference for honesty. For example, Fischbacher and Föllmi-Heusi (2013) found that 39% of participants were truthful in a private die-rolling experiment with guaranteed anonymity, Erat and Gneezy (2012) find evidence that individuals are reluctant to tell even Pareto white lies that benefit everyone, and Abeler et al. (2014) found that many participants in a representative German sample truthfully reported coin toss results over the phone, even though misreporting couldn't be detected and there were financial incentives to lie. Further evidence is provided by the influential paper by Abeler et al. (2019), which formalizes and tests various explanations for lying aversion using data combined from 90 different experimental studies. Their findings indicate that individuals not only prefer to be perceived as honest but also have an inherent preference for truth-telling.

This paper draws inspiration from and contributes to these bodies of literature in two key ways. First, it adds to the theoretical work by accounting for preferences for truth-telling in its framework, addressing a critical gap in the existing literature. Second, to the best of this author's knowledge, it is the first study to provide empirical evidence on how behavioral factors influence the expected consumer welfare gains from priority pricing (PP) schemes.

3 The Model

This section presents the theoretical framework, results, and corresponding hypotheses that guide the experimental analysis. The first part outlines the model setting and introduces an incentive-compatible priority pricing mechanism aimed at optimizing the allocation of two appointment slots based on agents' private valuations, from the perspective of a social planner seeking to maximize expected aggregate consumer welfare. The second part presents some theoretical results derived by comparing the expected aggregate consumer welfare under the priority pricing (PP) scheme to that of free-of-charge and uniform pricing schemes, while accounting for the presence of individuals with truth-telling preferences. The final part outlines the key experimental hypotheses designed to test the theoretical predictions.

3.1 Primitives

The model considers a simple setting in which two agents approach a service provider at the same time in order to secure an appointment. There are two available appointment slots: one on the same day and another on the following day. The service provider aims to allocate these appointment slots to maximize the expected aggregate utility of the agents.

Agents are expected utility maximizers, each characterized by a private type which determines their valuation for the appointment slot. The true type is denoted by the random variable θ, where $\theta \in \Theta = \{\theta_H, \theta_L\}$ and $\theta_H > \theta_L > 0$.[4] The types are drawn independently with $\mathbb{P}(\theta = \theta_H) = \alpha$ and $\mathbb{P}(\theta = \theta_L) = 1 - \alpha$, where $\alpha \in (0, 1)$.

Further, the agents' utility is decreasing in the amount of time they have to wait (t). The waiting time is determined by a function that depends on both the agent's reported type, denoted by $\hat{\theta}_i \in \Theta$, and the reported type of the other agent, $\hat{\theta}_j \in \Theta$. Formally, an agent i's utility function is defined as:

$$u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_j)) = \theta_i \delta^{t_i(\hat{\theta}_i, \hat{\theta}_j)}$$

where, $0 < \delta < 1$ represents the diminishing utility from a delayed appointment, $\hat{\theta}_i$ is the type that agent i reports, $\hat{\theta}_j$ is the reported type of other agent j and $t_i(\hat{\theta}_i, \hat{\theta}_j))$ is a function defined by:

$$t_i(\hat{\theta}_i, \hat{\theta}_j) = \begin{cases} 0 & \text{if } \hat{\theta}_i > \hat{\theta}_j, \\ 1 & \text{if } \hat{\theta}_i < \hat{\theta}_j, \\ X & \text{if } \hat{\theta}_i = \hat{\theta}_j, \end{cases}$$

where X is a random variable representing the allocation of the first appointment slot when $\hat{\theta}_i = \hat{\theta}_j$, with $\mathbb{P}(X = 1) = 0.5$ and $\mathbb{P}(X = 0) = 0.5$. This means that if agent i is the only one of the two agents to report as high type, agent i is assigned the first appointment slot (at $t = 0$). If agent i reports as low type while agent j reports as high type, then agent i is assigned the second slot (at $t = 1$). In the event that both agents report the same type, the first appointment slot is assigned randomly.

The agent knows their own type, θ_i, but does not know the type of the other agent. However, the distribution of types across agents is common knowledge. Therefore, agent i maximizes their expected utility by forming expectations over the possible types of the other agent j, based on this known distribution. Specifically, agent i seeks to maximize:

$$\mathbb{E}_{\theta_j}\left[u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_j))\right] = \mathbb{P}(\theta_j = \theta_H) u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_H)) + \mathbb{P}(\theta_j = \theta_L) u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_L))$$

Or, equivalently,

$$\mathbb{E}_{\theta_j}\left[u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_j))\right] = \alpha u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_H)) + (1 - \alpha) u(\theta_i, t_i(\hat{\theta}_i, \hat{\theta}_L))$$

[4] It is also assumed that $\theta_H > 2\theta_L$. This assumption is essential for ensuring the existence of the thresholds outlined in the subsequent propositions. Intuitively, this condition ensures that the difference in valuations between high-type and low-type agents is sufficiently large, which is necessary for the priority pricing scheme to generate the desired welfare outcomes.

Misreporting Incentives

In the absence of prices, in this set-up, high-type agents always maximize their expected utility by truthfully reporting their type, whereas low-type agents have a strategic incentive to misrepresent their type.[5]

If the other agent reports $\hat{\theta}_L$, the low-type agent's utility from truthful reporting is a weighted average of securing either the first or the second appointment slot, $u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_L)) = \frac{1}{2}\theta_L + \frac{1}{2}\theta_L\delta = \theta_L(\frac{1}{2} + \frac{1}{2}\delta)$, which is less than the utility from misreporting as $\hat{\theta}_H$, $u(\theta_L, t(\hat{\theta}_H, \hat{\theta}_L)) = \theta_L$, since $\delta \in (0, 1)$.

Similarly, if the other agents reports $\hat{\theta}_H$, the low-type agent's utility from truthful reporting is the utility derived from being appointed the second slot, $u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_H)) = \theta_L\delta$, which is less than the utility from misreporting as $\hat{\theta}_H$, $u(\theta_L, t(\hat{\theta}_H, \hat{\theta}_H)) = \frac{1}{2}\theta_L + \frac{1}{2}\theta_L\delta$.[6]

Therefore, without properly aligned pricing incentives, low-type agents will misreport their type to secure a better appointment slot, undermining the efficiency of the allocation process. Incentive-compatible pricing addresses this issue by aligning agents' incentives with truthful reporting.

Incentive-Compatible Priority Pricing

In the context of this simplified model, the optimal incentive-compatible pricing is derived by solving the following constrained optimization problem:

$$\max_{p_H, p_L} \ \mathbb{E}_{\theta_i, \theta_j} \left[\sum_{k \in \{i,j\}} \left(u(\theta_k) - p(\hat{\theta}_k) \right) \right]$$

$$\text{s.t.} \ \mathbb{E}_{\theta_j} \left[u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_j)) - p_L \right] \geq \mathbb{E}_{\theta_j} \left[u(\theta_L, t(\hat{\theta}_H, \hat{\theta}_j)) - p_H \right] \tag{IC-1}$$

$$\mathbb{E}_{\theta_j} \left[u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_j)) - p_H \right] \geq \mathbb{E}_{\theta_j} \left[u(\theta_H, t(\hat{\theta}_L, \hat{\theta}_j)) - p_L \right] \tag{IC-2}$$

$$\mathbb{E}_{\theta_j} \left[u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_j)) - p_L \right] \geq 0 \tag{IR-1}$$

$$\mathbb{E}_{\theta_j} \left[u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_j)) - p_H \right] \geq 0 \tag{IR-2}$$

where, $\hat{\theta}_k$ denotes the type reported by the agent, and $p(\hat{\theta}_k) = p_L$ if $\hat{\theta}_k = \hat{\theta}_L$ and $p(\hat{\theta}_k) = p_H$ if $\hat{\theta}_k = \hat{\theta}_H$. The incentive compatibility (IC) constraints ensure that agents are best off when they truthfully report their type conditional on the type of the other agent, and the individual rationality (IR) constraints ensures that participating is beneficial or at least not a loss, for every agent. The optimal pricing solution to this problem is obtained by setting $p_L^* = 0$ and $p_H^* = \frac{1}{2}\theta_L(1 - \delta)$, as detailed in Annex A.

[5]For high-type agents, truthful reporting always yields higher utility. When the other agent reports $\hat{\theta}_L$, the high-type agent's utility from truthful reporting, $u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_L)) = \theta_H$, exceeds the utility from misreporting, $u(\theta_H, t(\hat{\theta}_L, \hat{\theta}_L)) = \frac{1}{2}\theta_H + \frac{1}{2}\theta_H\delta$. Similarly, if the other agent reports $\hat{\theta}_H$, the utility from truthful reporting, $u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_H)) = \frac{1}{2}\theta_H + \frac{1}{2}\theta_H\delta$, exceeds the utility from misreporting, $u(\theta_H, t(\hat{\theta}_L, \hat{\theta}_H)) = \theta_H\delta$.

[6]This also applies to any uniform pricing scheme with a price that satisfies the individual rationality condition for low types, i.e., $\mathbb{E}_{\theta_j} \left[u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_j)) - p \right] \geq 0$, ensuring that low-type agents are never worse off by securing an appointment. In such cases, a constant price p would be subtracted from each utility calculation, which does not affect the comparison of utilities or the overall conclusion.

Incorporating Truth-Telling Preferences

To access how truth-telling preferences affect aggregate welfare under PP, the subsequent analysis differentiates low type agents based on their reporting behavior. Specifically, it is assumed that agent types are drawn from the set $\Theta' = \{\theta_H, \theta_{L_t}, \theta_{L_n}\}$.

Agents with the type θ_{L_t} have a valuation of θ_L for the appointment and are always truthful, i.e., they always report their true type or $\hat{\theta}_{L_t} = \hat{\theta}_L$. In contrast, agents with the type θ_{L_n}, despite having the same valuation θ_L, always misrepresent themselves as high type agents in the absence of properly aligned incentives, i.e., $\hat{\theta}_{L_n} = \hat{\theta}_H$.

Let α, β and γ denote the probabilities of an agent being high-type (θ_H), truthful low-type (θ_{L_t}), and non-truthful low-type (θ_{L_n}), respectively. These probabilities satisfy $\alpha+\beta+\gamma = 1$, with $0 < \alpha, \beta, \gamma < 1$, ensuring that each agent type occurs with positive probability. The distribution of these types is known to the social planner.

Expected Aggregate Consumer Welfare

Let $\mathbb{P}_{\theta_i\theta_j}$ denote the joint probability that the two agents arriving have the types θ_i and θ_j, respectively. The aggregate utility derived by the two agents from their allocated slots is given by:

$$U(\theta_i, \theta_j) = u(\theta_i, t(\hat{\theta}_i, \hat{\theta}_j)) + u(\theta_j, t(\hat{\theta}_j, \hat{\theta}_i)).$$

The total prices paid (or cost incurred) by the two agents is represented as:

$$P(\theta_i, \theta_j) = p(\hat{\theta}_i) + p(\hat{\theta}_j).$$

Therefore, the expected aggregate consumer welfare, considering both utilities and costs, is expressed as:

$$\mathbb{E}_{\theta_i,\theta_j}\left[U(\theta_i, \theta_j) - P(\theta_i, \theta_j)\right] = \sum_{\theta_i,\theta_j \in \Theta'} \mathbb{P}_{\theta_i\theta_j}(U(\theta_i, \theta_j) - P(\theta_i, \theta_j))$$

By linearity of expectation, this can be separated into expected aggregate utility minus the expected aggregate cost, as follows:

$$\mathbb{E}_{\theta_i,\theta_j}\left[U(\theta_i, \theta_j)\right] - \mathbb{E}_{\theta_i,\theta_j}\left[P(\theta_i, \theta_j)\right] = \sum_{\theta_i,\theta_j \in \Theta'} \mathbb{P}_{\theta_i\theta_j} U(\theta_i, \theta_j) - \sum_{\theta_i,\theta_j \in \Theta'} \mathbb{P}_{\theta_i\theta_j} P(\theta_i, \theta_j) \tag{1}$$

Under the optimal incentive-compatible PP scheme, only high type agents pay a price for the appointment ($p_H^* = \frac{1}{2}\theta_L(1-\delta)$), while low-type agents receive the service free of charge ($p_L^* = 0$). Thus, the expected cost per agent is $\mathbb{P}(\theta = \theta_H) \cdot p_H^* = \alpha p_H^*$ and for two agents this expectation is simply $2\alpha p_H^*$. Therefore, under incentive-compatible PP the expected aggregate consumer welfare simplifies to:

$$\mathbb{E}_{\theta_i,\theta_j}\left[U(\theta_i, \theta_j)\right] - \mathbb{E}_{\theta_i,\theta_j}\left[P(\theta_i, \theta_j)\right] = \sum_{\theta_i,\theta_j \in \Theta'} \mathbb{P}_{\theta_i\theta_j} U(\theta_i, \theta_j) - 2\alpha p_H^* \tag{2}$$

In the analysis that follows, the expected aggregate consumer welfare derived from the incentive-compatible PP scheme, where high type agents pay $p_H^* = \frac{1}{2}\theta_L(1-\delta)$ and low type agents receive the service free of charge, is compared to that from uniform pricing schemes, where the uniform price lies between zero and the expected cost under PP, $p \in [0, \alpha p_H^*]$.[7] Specifically, the analysis considers the following general case and two special cases:

1. **General Case:** The uniform price is set between zero and the expected cost under PP, $p \in (0, \alpha p_H^*)$. This scenario evaluates how intermediate pricing levels impact consumer welfare by comparing the expected aggregate welfare from PP, as given in equation (2), to $\sum_{\theta_i,\theta_j \in \Theta'} \mathbb{P}_{\theta_i\theta_j} U(\theta_i, \theta_j) - 2\alpha p_H^* \varepsilon$, for some $\varepsilon \in (0, 1)$ representing the relative pricing level.

2. **Bounding Case 1 (Free of Charge):** This case considers when no prices are charged, i.e., $p = 0$, corresponding to $\varepsilon = 0$ in the general case. It explores the welfare implications of PP relative to providing the service free of charge.

3. **Bounding Case 2 (Equal Expected Cost):** The uniform price is set equal to the expected cost under PP, i.e., $p = \alpha p_H^*$, which corresponds to $\varepsilon = 1$ in the general case. Here, the differences in expected costs across the pricing schemes are equalized, so the comparison focuses solely on the differences in the expected utilities arising from differences in the allocation of appointment slots.

3.2 Theoretical Results

By comparing the expected aggregate consumer welfare in equation (2) under the PP scheme with that of the three corresponding uniform pricing schemes, the following propositions are established. The proof for these propositions have been relegated to the Appendix B for brevity.

Proposition 1. *When the expected cost for agents under a uniform pricing scheme is less than under the incentive compatible priority pricing scheme, i.e. when the uniform price is set to $p' = \varepsilon \alpha p_H^*$, where $\varepsilon \in (0, 1)$, introducing incentive-compatible priority pricing (PP) will increase expected aggregate consumer welfare if and only if the probability of non-truthful, low-type agents (γ) exceeds the threshold $\gamma^\dagger$, where*

$$\gamma^\dagger = \frac{\theta_L(1-\varepsilon)}{\theta_H - \theta_L}.$$

Proposition 1 establishes that PP can enhance expected aggregate consumer welfare over any uniform pricing scheme with a lower expected cost, but only if the probability of non-truthful, low-type agents (γ) is sufficiently high, i.e. $\gamma > \gamma^\dagger$. This threshold $\gamma^\dagger$ decreases as the fraction ε increases, meaning that as the expected cost differential between PP and the uniform pricing narrows, a smaller probability of non-truthful, low-type agents is needed for PP to improve welfare.

[7] When the uniform price exceeds the expected cost under PP, i.e., $p > \alpha p_H^*$, PP will trivially generate higher expected aggregate consumer welfare by both increasing the expected utility from the allotted appointment slots and reducing the expected costs incurred by the agents.

Similarly, Proposition 2 addresses the first bounding case where $\varepsilon = 0$, meaning the expected cost differential between PP and the uniform pricing scheme is at its largest. In other words, this corresponds to the case where the service is provided free of charge to everyone under the uniform pricing scheme. It states that even in this case, implementing PP can enhance the expected aggregate consumer welfare if and only if the probability of non-truthful, low-type agents (γ) is sufficiently high, i.e. $\gamma > \gamma^*$.

Proposition 2. *When there is no cost to the agents, i.e. when the price is set to $p = 0$, introducing incentive-compatible priority pricing (PP) will increase expected aggregate consumer welfare if and only if the probability of non-truthful, low-type agents (γ) exceeds the threshold γ^*, where*

$$\gamma^* = \frac{\theta_L}{\theta_H - \theta_L}.$$

Finally, Proposition 3 addresses the second bounding case where $\varepsilon = 1$, meaning the expected cost between PP and the uniform pricing scheme are equal. In other words, this corresponds to the case where the uniform price $p = \alpha p_H^*$. The proposition states that the PP scheme always enhances expected aggregate consumer welfare compared to a uniform pricing scheme that imposes the same expected cost on agents. This result is driven by PP's theoretical ability to more accurately align appointment slot allocation with agents' true valuations. When the expected costs across the pricing schemes are identical, PP leads to higher aggregate consumer welfare by ensuring that the earlier appointment slot is allocated to the agent with the higher valuation, thereby, maximizing the expected aggregate utility.

Proposition 3. *When the expected cost for the agents under the incentive-compatible priority pricing (PP) scheme equals that under a uniform pricing scheme, i.e., when the uniform price is set to $p = \alpha p_H^*$, introducing priority pricing (PP) will always generate higher expected aggregate consumer welfare.*

3.3 Experimental Hypotheses

The implications of the theoretical findings presented in the preceding section are twofold: First, introducing priority pricing (PP) may not enhance expected aggregate consumer welfare unless the probability of non-truthful, low-need agents in the population is sufficiently large. Second, while PP could improve expected aggregate consumer welfare in populations with a low propensity for truthfulness, it could potentially reduce welfare in populations with a high propensity for truthfulness.

While the theoretical model focuses on aggregate welfare outcomes, budgetary constraints did not allow for a sample size with sufficient statistical power to measure these effects directly. However, the underlying mechanism presumed to influence welfare is the agents' tendency to truthfully disclose private information. Furthermore, the theoretical implications rest on the assumption that individuals will respond to the pricing incentive as predicted. Therefore, the experiment was designed to test the behavioral responses to priority pricing and its impact on populations with varying propensities for truth-telling. Specifically, the following hypotheses were tested:

Hypothesis 1. *The proportion of low-need individuals misreporting their type will be lower under priority pricing (PP) than when the service is offered free-of-charge, i.e.* $\pi_{fc} > \pi_{pp}$ *where pi denotes the proportion of non-truthful, low-need individuals, and the subscripts* fc *and pp denote free-of-charge and incentive-compatible priority pricing (PP), respectively.*

In other words, priority pricing has an effect in the desired direction and leads to a reduction in the proportion of low need individuals who misreport their level of need.

The next two hypotheses follow mechanically from the first one. They introduce a distinction between high-truthfulness and low-truthfulness populations, referring respectively to populations with higher and lower proportions of individuals with a propensity for truth-telling.

Hypothesis 2. *The proportion of low-need individuals who misreport their type will be lower under priority pricing (PP) than when the service is offered free-of-charge, across both high-truthfulness and low-truthfulness populations. Specifically,* $\pi_{fc_t} > \pi_{pp_t}$ *and* $\pi_{fc_n} > \pi_{pp_n}$*, where pi denotes the proportion of non-truthful, low-need individuals,* fc *denotes free-of-charge, pp denotes incentive-compatible priority pricing,* t *denotes the high-truthfulness population, and* n *denotes the low-truthfulness population.*

This is to say, the priority fee will have an effect in the desired direction regardless of the underlying propensity for truth-telling in populations.

Hypothesis 3. *The reduction in the proportion of non-truthful, low-need individuals under priority pricing (PP) will be greater for the high-truthfulness population than for the low-truthfulness population. This is denoted as:*

$$(\pi_{pp} - \pi_{fc})_t < (\pi_{pp} - \pi_{fc})_n$$

where, π *denotes the proportion of non-truthful, low-need individuals,* fc *denotes free-of-charge, pp represents incentive-compatible priority pricing,* t *denotes the high-truthfulness population, and* n *denotes the low-truthfulness population.*

Given that the low-truthfulness population has a higher proportion of individuals likely to be non-truthful in the absence of incentives compared to the high-truthfulness population, priority pricing should have a greater effect on reducing non-truthful behavior in the low-truthfulness group.

4 Experimental Design

To test the hypotheses, a between-subjects online experiment was conducted using a survey-based approach. The experiment was pre-registered in the AEA RCT Registry (Thami, 2024). The survey was programmed in Qualtrics, and the participants were recruited through Prolific. A total of 696 participants completed the survey, which took an average of approximately 12 minutes to complete.[8] Participants earned an average of \$5, comprising a \$2 participation fee in accordance with the wage rate recommended by Prolific, and an additional average of \$3 in experimental earnings.

[8]Although 700 participants completed the survey, four submissions were excluded from the study based on recommendations from Prolific due to concerns over shared IP addresses, which violates Prolific's policy.

4.1 Online Experiment

The experimental design closely resembles the one used in Campos-Mercade (2022), where participants' types are first observed and later successfully used to vary the composition of types in sub-groups. In this study, participants' propensity for truth-telling is first observed and then used in two ways: first, to shift the participants' beliefs about the truth-telling propensities of others, encouraging responses consistent with these beliefs; and second, to conduct a weighted analysis simulating high- and low-truthfulness populations.

The survey used to collect data for the experiment consisted of three sections, requiring participants to make six decisions about whether to truthful disclose private information. Participants were randomly assigned to one of the following four conditions: 1) Free of Charge with Truthful Prompt; 2) Free of Charge with Non-Truthful Prompt; 3) Priority Pricing with Truthful Prompt; or 4) Priority Pricing with Non-Truthful Prompt. Table 1 presents the distribution of participants across these conditions.

Table 1: Distribution of Participants Across Experimental Conditions (%)

	Free of Charge	Priority Pricing	**Total**
Truthful Prompt	50.69	49.25	50.00
Non-Truthful Prompt	49.31	50.75	50.00
Total	52.16	47.84	100.00

In each section, participants were first introduced to the setup of the interaction, the two potential types they could be assigned as, and the associated payoffs for each type. They were then randomly assigned a type and had to decide whether to disclose their assigned type truthfully. Participants made decisions as both types, with the order of type allocation randomized. Figure 1 details the differences in prompts across sections according to the assigned condition.

Furthermore, at the beginning of the survey, participants were informed that they would receive additional compensation based on the outcome of one randomly selected decision from the first two sections. They were also notified of the possibility to earn an extra bonus from their decisions in Section 3. This ensured that each decision carried the potential for increased earnings.

Section 1 sought to identify participants with a higher propensity for being non-truthful. While common identification measures typically rely on the use of experimental paradigms such as sender-receiver games (Gneezy, 2005),[9] die-roll tasks (Fischbacher & Föllmi-Heusi,

[9] In sender-receiver games, one participant (the sender) knows the payoffs of different options and chooses to send a truthful or false message to a second participant (the receiver), whose decision determines the final payoffs for both players, creating a dilemma for the sender between truth-telling and maximizing personal gain.

Figure 1: Experimental Design

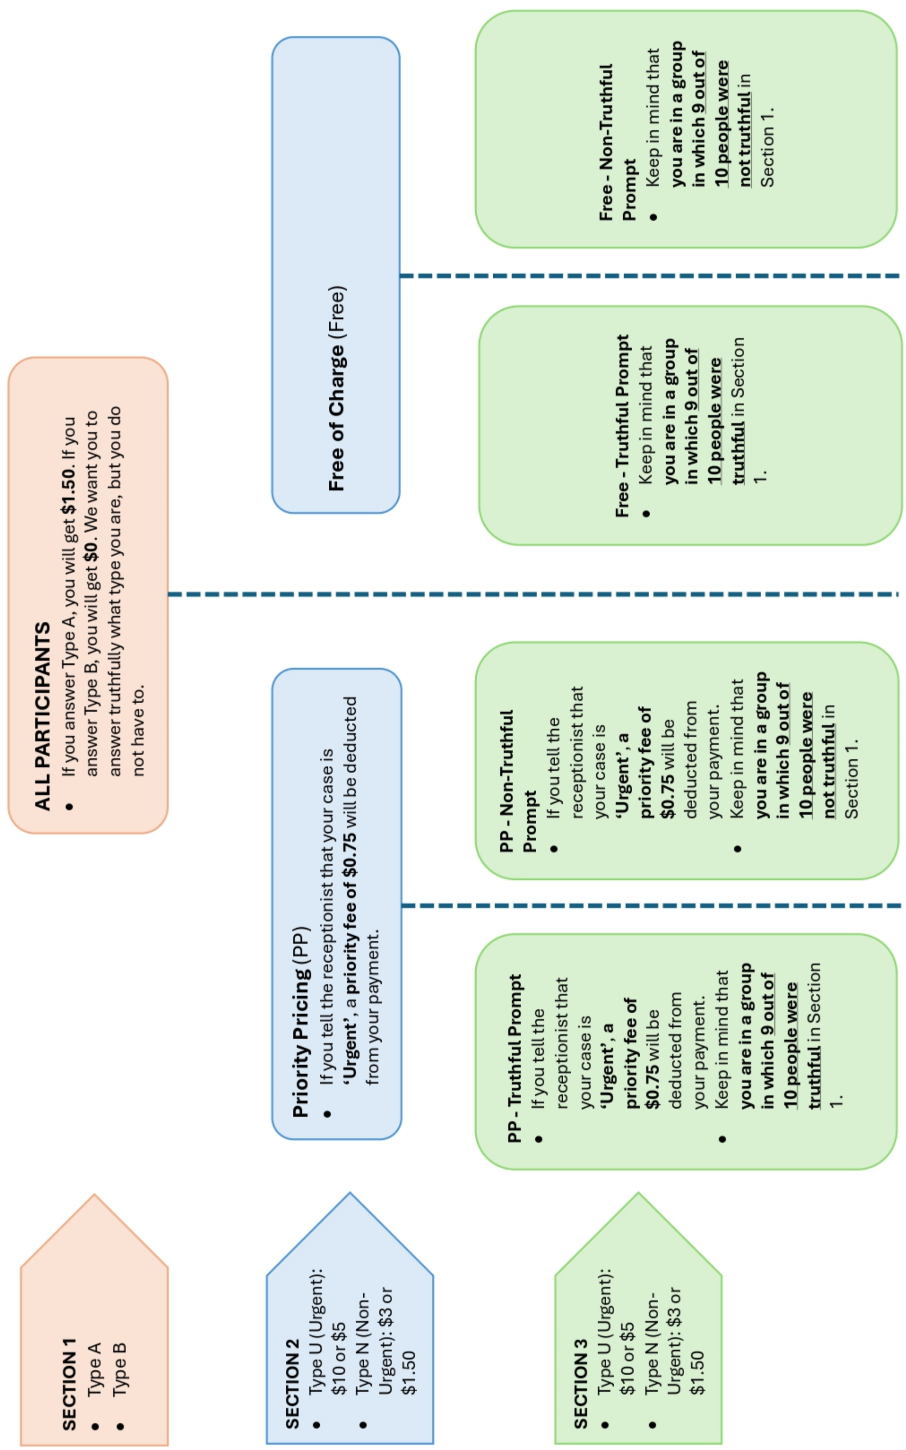

2013),[10] coin flip tasks (Bucciol & Piovesan, 2011),[11] and matrix tasks (Mazar et al., 2008),[12], here a more direct approach was adopted. Participants were informed that they would randomly be assigned as either a Type A or Type B individual. They were told that they would receive \$1.50 for saying that they are Type A and \$0 for Type B, irrespective of their actual type assignment. They then responded to the question "What is your Type?" This approach offered two key advantages over other established paradigms: 1) it allowed for the identification of non-truthful responses at the individual level, enabling the creation of weights to simulate high- and low- truthfulness populations; and 2) the structure of this interaction and the size of the monetary incentive closely mirrors those in the subsequent sections, which are central to the experiment.

Section 2 captured the baseline effect of adopting priority pricing on truth-telling behavior, compared to free-of-charge service provision, providing the data to test Hypothesis 1. In it, participants simulated the scenario outlined in the theoretical model. They were informed that they would be randomly paired with another participant to simulate a queuing scenario for booking a doctor's appointment with two available slots: one immediate and one for the following day.

The two potential types were Type U (Urgent), with a higher valuation for the appointment slots (\$10 for the immediate one and \$5 for the next day), and Type N (Non-Urgent), with a lower valuation for the appointment slot (\$3 for the immediate and \$1.50 for the next day). To address potential concerns about inequality aversion, Type N participants were guaranteed an additional payment of \$7 to ensure perceived fairness in compensation across types. This was particularly important given that the primary payment was determined by a randomly selected decision.

In keeping with the theoretical model's assumption of known type distribution, participants were informed that Type U would be assigned with a 25% probability and Type N with a 75% probability. Further, participants in the priority pricing treatment groups were also informed that a priority fee of \$0.75 would be applied if they answered that they were Type U (Urgent).[13] They were also explicitly informed that this fee was set such that reporting their type truthfully would maximize their expected payoff.

Participants were also briefed on the appointment allocation rules, which were identical to those outlined in the theoretical model. If one participant reported Type U while the other reported Type N, the Type U participant would receive the first appointment. If both participants reported the same type, the first appointment would be randomly assigned to

[10] In die-roll tasks, participants report a private roll outcome linked to a payment, enabling researchers to estimate dishonesty by comparing reported averages to expected random distributions.

[11] In coin-flip tasks, participants report the outcome of a self-generated coin toss, with rewards for one outcome, allowing researchers to estimate dishonesty by comparing reported win rates to the expected random 50/50 distribution, though individual dishonesty cannot be directly detected.

[12] In matrix tasks, participants search for number pairs in matrices that add up to 10.00 and self-report their performance for a monetary reward, with dishonesty measured by comparing reported versus verified solutions, either at the aggregate level by contrasting self-graded and experimenter-graded groups or at the individual level through direct self-reporting.

[13] Recall that in the optimal incentive-compatible priority pricing scheme, low-type agents receive the service free of charge, while high-type agents pay a fee of $p_H^* = \frac{1}{2}\theta_L(1-\delta)$. The priority fee was calculated assuming a discount factor, δ of 0.5. Given, that in this setup $\theta_L = \$3$, this results in a priority fee of $p_H^* = \frac{3}{2}(1-0.5) = \$0.75$.

one of them.

In **Section 3**, participants were presented with the same scenario and made the same decision as in Section 2, with the only difference being that, prior to making the decision, they were informed that they would be randomly assigned to a group of 10 participants and could qualify for additional payment based on their group's characteristics. Specifically, those assigned to the Free-Truthful and PP-Truthful Prompt conditions were told that they would receive extra payment from the outcome of their decision in the section only if 9 out of 10 group members had been truthful in Section 1. Conversely, those assigned to the Free-Non-Truthful and PP-Non-Truthful Prompt conditions were told that they would receive extra payment only if 9 out of 10 had been non-truthful in Section 1. This section provided the data necessary to test Hypotheses 2 and 3.

By introducing this additional payment criteria based on group characteristic, the design created an incentive for participants to respond as if their group composition aligned with the specified criteria. For instance, participants assigned to the Free-Truthful Prompt and PP-Truthful Prompt conditions were encouraged to respond as they might if most of the others in their group were likely to be truthful. Conversely, those assigned to the Non-Truthful prompt conditions were encouraged to respond as they might if most of the others were likely to be non-truthful.

This experimental design enabled the simulation of high- and low-truthfulness populations in two ways. First, it allowed for the assignment of probability weights to participant responses based on whether they had been truthful or non-truthful when assigned as Type B in Section 1.[14] Specifically, in the Truthful Prompt conditions, responses from participants who were truthful in Section 1 were assigned a probability weight of 1.8, while responses from non-truthful participants were weighted at 0.2 to simulate a high-truthfulness population. Conversely, in the Non-Truthful Prompt conditions, responses from participants who had been non-truthful in Section 1 were weighted at 1.8, and those from participants who had been truthful were weighted at 0.2 to simulate a low-truthfulness population. Second, the additional payment criteria introduced in Section 3 encouraged participants to respond as they would in a group predominantly composed of truthful or non-truthful individuals, further aligning participant responses with the simulated population characteristics.

Additionally, since the prompt regarding the criteria for additional payment was the only source of variation between the two Priority Pricing (PP) conditions and, separately, the two Free-of-Charge (Free) conditions, this design also enables the study of how beliefs about others' propensities for truth-telling influence the effect of priority pricing.

Furthermore, to ensure that participants fully understood the experimental setup, they

[14] Probability weights are factors assigned to observations to adjust their influence in an analysis. Probability weights were calculated as the inverse of the relative selection probability given the observed proportion in the sample and an assumed underlying population distribution. For example, the probability weight for participants who had been truthful in Section 1 and were assigned to one of the Truthful Prompt conditions was derived as follows. Assuming the participant was randomly drawn from a high-truthfulness population—one where 90% are likely to be truthful and only 10% non-truthful—their selection probability in this theoretical population would be 0.9. However, in the experiment, there were 175 participants who were truthful in Section 1 of a total of 348 participants assigned to the Truthful Prompt conditions giving a sample selection probability of $\frac{175}{348} \approx 0.5$. Since the assumed underlying population has 90% truthful individuals but the actual sample only had 50%, the selection probability for truthful individuals relative to the population distribution is $\frac{0.5}{0.9}$. The probability weight is then determined by taking its inverse, i.e., $\frac{0.9}{0.5} = 1.8$.

were required to successfully answer a series of comprehensive control questions before being allowed to advance to the main decision questions.[15] These questions were designed to ensure a thorough understanding of the experimental setup and the potential payoffs. Participants in the treatment groups answered additional questions about the priority fee, which clarified that truthful responses would maximize their potential earnings. Participants were allowed to reattempt the control questions as many times as needed to ensure they thoroughly understood the experimental setup.

5 Results

This section presents descriptive results and the test results for the hypotheses outlined in Section 3.3. Additionally, it also includes results from simulation exercises that demonstrate the impact of observed behaviors on aggregate consumer welfare.

5.1 Descriptive Results

Table 2 presents the distribution of reported type by assigned type in Section 1. Nearly all participants (98.85% or 688 participants) reported their type truthfully when assigned as Type A. In contrast, only about half (49.84% or 347 participants) reported their type truthfully when assigned as Type B. This response pattern is consistent with the preference for truth-telling that motivates this study: participants are truthful when incentives are aligned, and a substantial proportion are truthfully even when incentives are misaligned.

Table 2: Distribution of Reported vs. Assigned Type in Section 1

		Assigned Type	
		Type A	**Type B**
Reported Type	**Type A**	98.85%	50.14%
	Type B	1.15%	49.84%

Table 3 presents the distribution of responses in Section 2. The vast majority of participants responded truthfully when assigned as Type U (Urgent), with 96.42% of participants in the Free of Charge condition reporting their type truthfully and 93.09% under Priority Pricing. However, a significant proportion of participants misreported their level of need when assigned as Type N (Non-Urgent), especially when appointments were offered free of charge. Specifically, nearly half of Type N participants (49.04%) reported as Type U (Urgent) when appointments were free, compared to only 26.73% when a priority fee was charged.

Table 4 presents the distribution of responses in Section 3 under the Truthful and Non-Truthful prompt conditions. The general pattern is similar to the one observed in Section 2 (Table 3), with a high proportion of truthful responses from participants when assigned as Type U (Urgent). Additionally, when participants were assigned as Type N (Non-Urgent),

[15]While this meant there was a high attrition rate of approximately 18%, this measure was crucial to ensuring participant attention and comprehension.

Table 3: Distribution of Reported vs. Assigned Type by Pricing Condition in Section 2

		Assigned Type			
		Free of Charge		**Priority Pricing**	
		Non-Urgent	**Urgent**	**Non-Urgent**	**Urgent**
Reported Type	**Non-Urgent**	50.96%	3.58%	73.27%	6.91%
	Urgent	49.04%	96.42%	26.73%	93.09%

the proportion of non-truthful responses is notably lower under priority pricing compared to the free-of-charge condition, across participants in both the Truthful and Non-truthful prompt conditions. However, there are notable differences across the two prompt conditions that merit further discussion.

In the Truthful Prompt condition, a substantial increase in truthful reporting is observed for Type N (Non-Urgent) participants compared to the baseline case in Section 2, suggesting that when participants believed others would likely report truthfully, they were more inclined to do the same. Specifically, in the Free of Charge condition, 60.34% of the participants reported truthfully when assigned as Type N (compared to 50.96% in Section 2), and in the Priority Pricing condition, truthful reporting increased to 82.25% (up from 73.27% in Section 2).

Conversely, under the Non-Truthful Prompt condition, the opposite pattern is observed, with there being an increase in misreporting relative to Section 2. In the Priority Pricing condition, only 64.63% of Type N participants reported truthfully, a noticeable drop from the Truthful Prompt condition (82.25%) and from Section 2 (73.27%). This suggests that when participants believed others were more likely to misreport truthfully, they were more inclined to misreport under Priority Pricing. However, the Free of Charge condition a small increase in truthful reporting, with 53.80% reporting truthfully compared to 50.96% in Section 2.

5.2 Main Results

The analysis presented in this section was pre-registered in the AEA RCT Registry Thami (2024). The sample size in this study provided sufficient power to detect the main effect.[16]

Figure 2 shows the proportion of participants who were non-truthful about their type when assigned as Type N (Non-Urgent) in Section 2 of the survey. Hypothesis 1 posits that the priority pricing should have a statistically significantly lower proportion of non-truthful, low-need participants compared to the free-of-charge control groups. A Wilcoxon

[16]The power analysis was conducted using a simulation-based approach (STATA code available upon request) with a latent variable model assuming a standard normal distribution for truth-telling propensity. Key assumptions included a 55% explanatory power of non-truthfulness in Section 1 on later behavior, a baseline non-truthful behavior of approximately 50% in the absence of a priority fee, and a priority fee effect size of 56%. This analysis indicated 100% power to detect the main effect (Hypothesis 1), 96.4% power for detecting the treatment effect in the Truthful prompt conditions and 100% for detecting it in the Non-Truthful prompts conditions, and 87.3% power for detecting an interaction effects between priority pricing and truthful prompts conditions at the 5% significance level.

Table 4: Distribution of Responses by Assigned Type and Pricing Condition for Truthful and Non-Truthful Prompt Conditions in Section 3

(a) Truthful Prompt

		Assigned Type			
		Free of Charge		**Priority Pricing**	
		Non-Urgent	**Urgent**	**Non-Urgent**	**Urgent**
Reported Type	**Non-Urgent**	60.34%	3.35%	82.25%	4.73%
	Urgent	39.66%	96.65%	17.75%	95.27%

(b) Non-Truthful Prompt

		Assigned Type			
		Free of Charge		**Priority Pricing**	
		Non-Urgent	**Urgent**	**Non-Urgent**	**Urgent**
Reported Type	**Non-Urgent**	53.80%	7.07%	64.63%	12.80%
	Urgent	46.20%	92.93%	35.37%	87.20%

rank-sum test strongly supports this hypothesis, showing a significant difference ($p < 0.01$) between the priority pricing conditions (N = 333) and the free-of-charge conditions (N = 363). Specifically, 49% of participants in the free-of-charge control groups were non-truthful, compared to only 26.7% of those in the priority pricing treatments, representing a 45.5% reduction in non-truthful responses.

Result 1. *The proportion of non-truthful, low-type participants, i.e., Type N (Non-Urgent) reporting as Type U (Urgent), is statistically significantly lower under priority pricing compared to the free of charge condition.*

To test Hypothesis 2, samples representing high- and low-truthfulness populations were generated based on participants' responses in Section 1 when assigned as Type B. For the high-truthfulness groups, only observations from participants in the Truthful prompt conditions were included—specifically, from participants in the PP-Truthful prompt condition for the priority pricing group and from the Free-Truthful prompt condition for the free-of-charge group. Each sample was constructed to ensure that at least 90% of participants had truthfully reported as Type B in Section 1.

Similarly, the low-truthfulness groups were created using only observations from participants assigned to the Non-Truthful prompt conditions. Observations for the priority pricing group came from the PP-Non-Truthful prompt condition, while those for the free-of-charge group were drawn from the Free-Non-Truthful prompt condition. These samples were constructed to ensure that at least 90% of participants had misreported as Type A when assigned as Type B in Section 1.

Figure 2: Proportion of Participants who were Non-Truthful (Section 2)

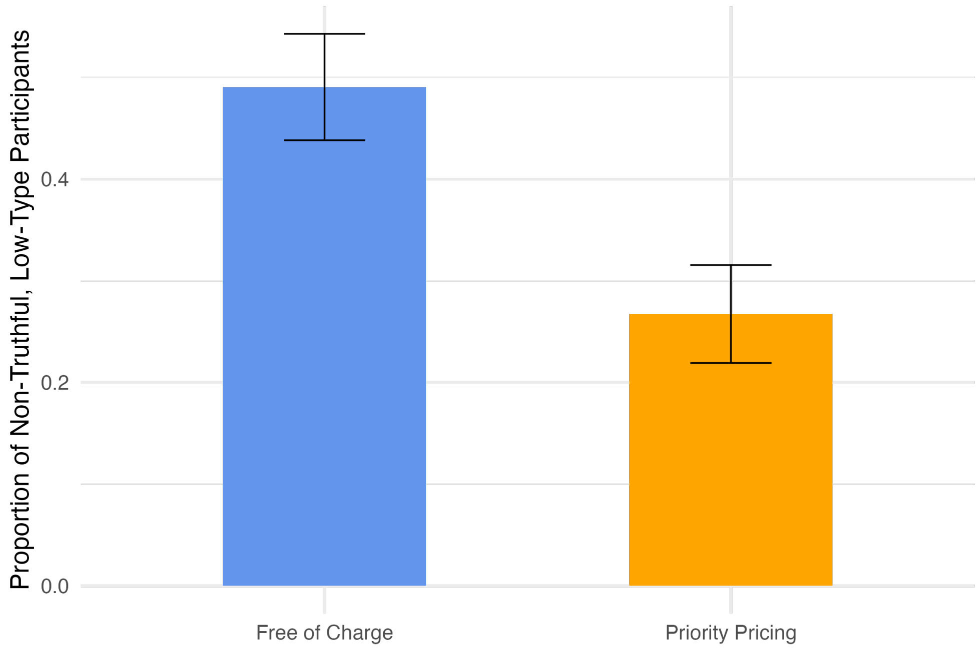

Note: The confidence intervals were computed using bootstrap resampling with 10,000 iterations.

To account for potential variability in the selection process, 10,000 samples were generated following the criteria above, and the Dunn test was conducted on each sample.[17] The primary variable of interest was participants' responses as Type N (Non-Urgent) in Section 3 of the survey.

Hypothesis 2 states that the proportion of non-truthful, low-need participants should be lower under priority pricing compared to the free of charge condition, across both high- and low-truthfulness populations. Figure 3 presents the average proportion of non-truthful, low-need participants (i.e., participants who misreported when assigned as Type N in Section 3) across 10,000 samples under Free of Charge and Priority Pricing conditions for both high-truthfulness and low-truthfulness groups.

Dunn test consistently rejects the null hypothesis that the proportion of non-truthful, low-need participants are equal across all four groups, with $p < 0.01$ in all samples in each of the 10,000 bootstrap samples. Furthermore, the tests reveal a statistically significant difference between the priority pricing and free of charge conditions within the low-truthfulness groups, with $p < 0.05$ observed in 94.8% of the samples. Conversely, no statistically significant difference was found within the high-truthfulness groups. However, it is notable that the bootstrapped confidence intervals, constructed by taking the 2.5^{th} and 97.5^{th} percentile observations from the 10,000 samples, do not overlap between the priority pricing and free of charge conditions for either the high-or low-truthfulness groups. This non-overlap suggests

[17]The tests were conducting using STATA's dunntest command with Bonferroni adjustment to account for multiple comparisons.

Figure 3: Proportion of Participants who were Non-Truthful (Section 3)

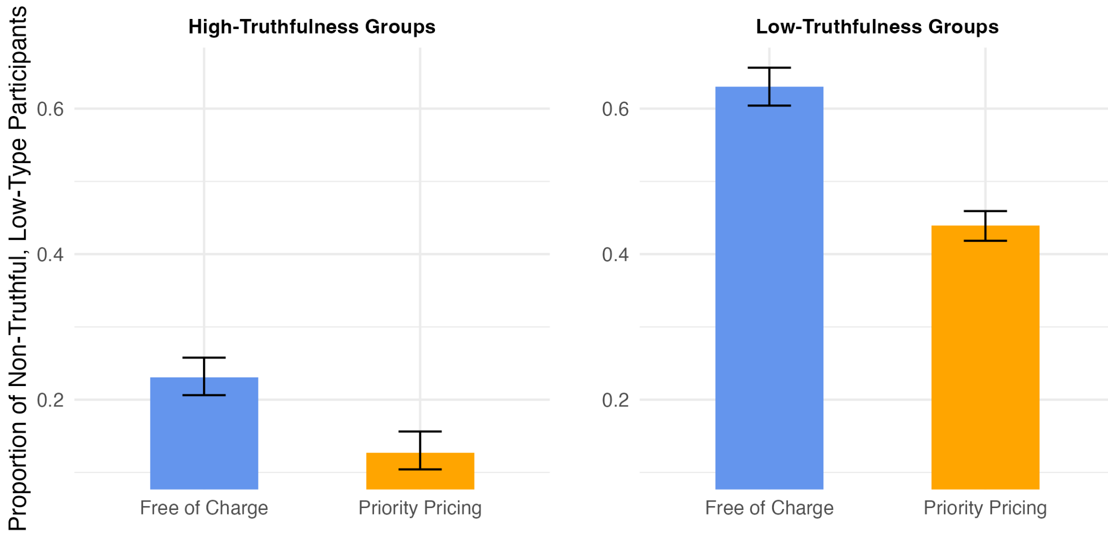

Note: The confidence intervals were calculated by using the 2.5^{th} and 97.5^{th} percentiles of the proportion of non-truthful, low-type participants from the 10,000 subsamples.

a consistent directional difference, even in the absence of statistically significant results in the high-truthfulness group.

Result 2. *The proportion of non-truthful, low-need participants, i.e. Type N (Non-Urgent) reporting as Type U (Urgent), is significantly lower under priority pricing condition compared to the free of charge condition within the low-truthfulness groups.*

Hypothesis 3 posits that priority pricing treatment will have a larger effect in reducing non-truthful behavior in low-truthfulness populations compared to high-truthfulness populations. To test this hypothesis, probability-weighted regressions are conducted. The probability weights simulate the desired distribution of truthful and non-truthful individuals within each experimental condition, thus enabling comparisons across high- and low-truthfulness populations. In the results table, the truthful prompt assignment is labeled as the high-truthfulness group, reflecting the probability-weighted approach to simulate the high- and low-truthfulness distributions.

The primary variable of interest is a binary indicator for whether the participant misreported when assigned as Type N in Section 3. The key coefficient of interest is the interaction between assignment to the Priority Pricing condition and the Truthful prompt condition which captures the differential effect of priority pricing across high-truthfulness and low-truthfulness groups.

Table 5 presents the results. Column (1) indicates that while priority pricing appears to reduce non-truthful reporting, the interaction effect between belonging to a high-truthfulness group and assignment to the priority fee condition, contrary to Hypothesis 3, is not statistically significant. Columns (3) and (5) present the logitistic and probit regression results, which corroborate the direction, effect size and statistical significance observed in the OLS

Table 5: Differential Treatment Effect of Priority Pricing Across High- and Low-Truthfulness Groups (Probability Weighted Analysis)

	OLS		Logit		Probit	
	(1)	(2)	(3)	(4)	(5)	(6)
Priority Price	–0.183***	–0.178***	–0.743***	–0.735***	–0.464***	–0.455***
	(0.067)	(0.067)	(0.276)	(0.283)	(0.172)	(0.174)
High-Truthfulness Group	–0.388***	–0.376***	–1.683***	–1.657***	–1.036***	–1.018***
	(0.060)	(0.060)	(0.289)	(0.293)	(0.173)	(0.175)
Priority Price X High-Truthfulness Group	0.075	0.079	–0.001	0.016	0.047	0.045
	(0.083)	(0.083)	(0.449)	(0.453)	(0.260)	(0.261)
Constant	0.623***	0.738***	0.502***	1.158**	0.313***	0.716***
	(0.046)	(0.087)	(0.194)	(0.460)	(0.120)	(0.271)
Controls	No	Yes	No	Yes	No	Yes
Observations	696	696	696	696	696	696
R-squared	0.158	0.174				

[a] This table reports the results of OLS, logitistic, and probit regressions that test for the differential effect of priority pricing across truthful and non-truthful prompt conditions. The binary outcome variable takes the value of 0 if the participant truthfully reported their type as Type N when assigned as Type N in Section 3 and 1 otherwise. Priority Price is a dummy variable indicating if the participant had been assigned to one of the priority price conditions. High-Truthfulness Group is a dummy variable indicating if the participant belonged to one of the Truthful prompt conditions. Priority Price × High-Truthfulness Group captures the interaction effect of being assigned to both the priority fee condition and the truthful prompt condition, and it is the primary variable of interest in this analysis. Participant age, sex and employment status are used are control variables.
[b] Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

regression. Columns (2), (4) and (6) presents the OLS, logistic and probit regression estimates, respectively, including control variables sex, age and employment status.[18] These control variables were included because they were available as part of the demographic data automatically provided by Prolific, and there is substantial evidence indicating their relevance in the context of truth-telling (see, for example, Gerlach et al., 2019).[19] The results across all models are consistent.

Result 3. *Contrary to Hypothesis 3, the analysis reveals no statistically significant difference in the impact of priority pricing between the high-truthfulness and low-truthfulness groups.*

5.3 Role of Beliefs in the Experimental Outcome

While the predicted effect is not observed in the weighted regressions, a significant nonlinear interaction effect emerges in the non-weighted regressions (Table 6, Columns (3) to (6)). The only source of systematic variation across the Truthful and Non-Truthful prompt conditions was the prompt itself, i.e., whether participants were informed that qualifying for additional

[18] Employment status is a categorical variable that classifies each participant's employment status as either "Employed," "Unemployed," "Other," or "Missing.".

[19] Research suggests that age and sex may be correlated to truth-telling behavior. Specifically, studies have found that men and younger individuals may have a higher propensity for non-truthfulness. Furthermore, employment status is included as a control variable because it may be associated with systematic differences in participants' financial motivations.

payment required 9 out of 10 participants in their group to be either truthful or non-truthful in their responses in Section 1.

These results suggest that participants adjust their truth-telling behavior based on their beliefs about the truthfulness of others in their group. Specifically, the significant negative interaction effect between being assigned to a truthful prompt condition and the priority pricing condition indicates that priority pricing significantly reduces the likelihood of non-truthful reporting when participants believe others are more likely to be truthful, compared to when they believe others are less likely to be truthful.

Hypothesis 3 relies on the idea that because there are more non-truthful participants in the low-truthfulness groups, the introduction of priority pricing should have a more pronounced effect in reducing non-truthful behavior within these groups compared to high-truthfulness groups. However, the tendency of participants to adjust their behavior based on their beliefs about the truthfulness of others may be weakening the expected distributional effects, which could potentially explain the lack of significant results in the weighted regressions.

Table 6: Differential Treatment Effect of Priority Pricing Across Truthful and Non-Truthful Groups (Non-Weighted Analysis)

	OLS		Logit		Probit	
	(1)	(2)	(3)	(4)	(5)	(6)
Priority Price	–0.108**	–0.111**	–0.451**	–0.474**	–0.280**	–0.295**
	(0.053)	(0.052)	(0.220)	(0.224)	(0.137)	(0.138)
Truthful Prompt	–0.065	–0.068	–0.267	–0.287	–0.167	–0.179
	(0.052)	(0.052)	(0.213)	(0.217)	(0.133)	(0.134)
Priority Price X Truthful Prompt	–0.111	–0.097	–0.663**	–0.607*	–0.383*	–0.348*
	(0.071)	(0.071)	(0.336)	(0.339)	(0.201)	(0.203)
Constant	0.462***	0.556***	–0.152	0.302	–0.096	0.181
	(0.037)	(0.076)	(0.148)	(0.349)	(0.093)	(0.211)
Controls	No	Yes	No	Yes	No	Yes
Observations	696	696	696	696	696	696
R-squared	0.049	0.064				

[a] This table reports the results of OLS, logitistic, and probit regressions that test for the differential effect of priority pricing across truthful and non-truthful prompt conditions. The binary outcome variable takes the value of 0 if the participant truthfully reported their type as Type N when assigned as Type N in Section 3 and 1 otherwise. Priority Price is a dummy variable indicating if the participant belonged to one of the priority price conditions. Truthful prompt is a dummy variable indicating if the participant belonged to one of the truthful prompt conditions. Priority Price × Truthful Prompt captures the interaction effect of being assigned to both the priority price and the truthful prompt condition, and it is the primary variable of interest in this analysis. Participant age, sex and employment status are used as control variables.
[b] Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

5.4 Implications for Expected Aggregate Welfare

The primary aim of this section is to demonstrate how the observed behavioral patterns in the experimental data affects aggregate welfare outcomes. To this end, this section presents results from three different simulation exercises, each of which calculates and compares the average aggregate payoffs across different pricing schemes using the experimental data and evaluates these outcomes against their theoretical expectations. In these simulations, 10 participants were sampled with replacement 50,000 times from each of the four experimental conditions. In line with the setup in the theoretical model, participants within each sample were randomly paired. Each participant was assigned either Type U (Urgent) with a probability of 0.25 or Type N (Non-Urgent) with a probability of 0.75. These assigned types, along with the participants' actual responses corresponding to their assigned types, were used to determine the appointment allocations and payoffs for each pair. The average aggregate payoff was then calculated for each sample. Although the underlying sample size lacks sufficient power to detect statistically significant differences, the simulations nonetheless provide valuable insights.

5.4.1 Baseline Comparison of Aggregate Payoff Across Pricing Schemes

Figure 4: Average Aggregate Payoff in Section 2

Figure 4 presents the average aggregate payoff calculated using data from Section 2, along with the theoretical expectation, with results shown for each pricing scheme: priority pricing (using data from the priority pricing condition), and free-of-charge and a uniform pricing scheme (using data from the free-of-charge condition). Under th uniform pricing scheme, all participants are charged a uniform price of \$0.1875, calibrated so that the total expected cost for each pair of participants under the priority pricing scheme equals the expected cost under the uniform pricing scheme.[20]

According to theoretical expectations, the average aggregate payoff from priority pricing will exceed that from the free of charge condition if and only if the probability of low-need, non-truthful participants exceeds the threshold, $\gamma^* = \frac{\theta_L}{\theta_H - \theta_L}$ (which is 0.429 in this experimental setting).[21] However, the empirical probability of non-truthful Type N participants

[20]The uniform price is set to $p_{\text{uniform}} = \alpha p_H^*$. Given that α, the probability of being assigned Type U, is 0.25, and $p_H^* = 0.75$, the uniform price is calculated as $p_{\text{uniform}} = 0.25 \times 0.75 = \$0.1875$.

[21]Given, $\theta_H = 10$ and $\theta_L = 3$, the threshold γ^* is calculated as $\gamma^* = \frac{3}{10-3} = 0.429$.

in the free of charge conditions, was 0.367, which falls below this threshold.[22] Consequently, the theoretical expected payoff for the free of charge scenario was \$7.46 which exceeded the \$7.41 expected for priority pricing. The expected payoff for uniform pricing was \$7.09, the lowest among the three.

Moreover, contrary to theoretical predictions, the empirical average aggregate payoff from uniform pricing was higher than that from priority pricing. Priority pricing resulted in the lowest average aggregate payoff among the three pricing schemes, showing the greatest deviation from theoretical expectation. This result is driven by the fact that not all participants responded to the incentive under priority pricing - i.e., a substantial proprotion of participants misreported even under priority pricing (26.7%). These findings highlight that not all individuals respond to the incentives under priority pricing as anticipated, suggesting that the actual threshold for priority pricing to be welfare enhancing could be substantially higher than theoretical predictions suggest.

5.4.2 Population Truthfulness Propensity and Average Aggregate Payoffs Across Pricing Schemes

Figure 5 presents the average aggregate payoff from Section 3 across the three pricing schemes, comparing high-truthfulness and low-truthfulness groups. The figure also includes the theoretical expected payoff. Samples were selected based on responses from Section 1. Each high-truthfulness group consisted of 10 participants drawn from the truthful prompts condition within each pricing scheme (free-of-charge and priority pricing): nine participants were randomly selected from those who had truthfully identified as Type B when assigned as such in Section 1, and one was randomly selected from those who had been non-truthful. Conversely, each low-truthfulness group was drawn from the non-truthful prompts condition within each pricing scheme and included nine participants randomly selected from those who had misrepresented themselves as Type A when assigned as Type B, along with one randomly selected truthful participant.

The empirical probability of non-truthful Type N participants in the free of charge conditions were 0.176 in the high-truthfulness groups and 0.47 in the low-truthfulness groups. Therefore, according to theoretical predictions, the average aggregate payoff from priority pricing should exceed that of the free-of-charge case for the low-truthfulness groups, while it should be lower for the high-truthfulness groups. As shown in Figure 5, the latter expectation aligns with the experimental data, but the former is contradicted. Specifically, priority pricing is found to yield lower average aggregate payoffs across both the high- and low-truthfulness groups.

This is because priority pricing did not completely eliminate misreporting and it was also less effective at deterring misrepresentation among participants in the low-truthfulness groups, as discussed in the preceding sections. To add further clarity, the average proportion of non-truthful Type N participants in the priority pricing conditions were 0.097 in the high-truthfulness groups, compared to 0.326 in the low-truthfulness groups. This is to say that while priority pricing reduced non-truthful reporting by about 44.89% in the high-truthfulness group, it only reduced it by about 30.64% in the low-truthfulness group.

[22] The empirical probability is calculated as the proportion of observed non-truthful Type N participants in the generated samples.

Figure 5: Average Aggregate Payoff in Section 3

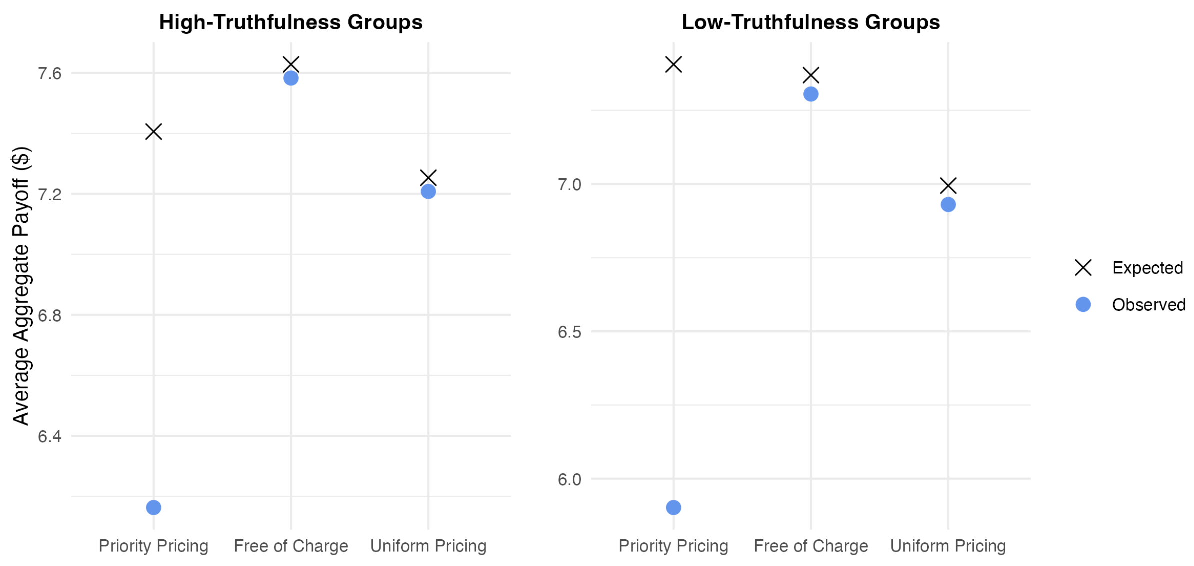

5.4.3 Beliefs and Average Aggregate Payoffs Across Pricing Schemes

Figure 6: Comparison of Average Aggregate Payoff in Section 2 and Section 3

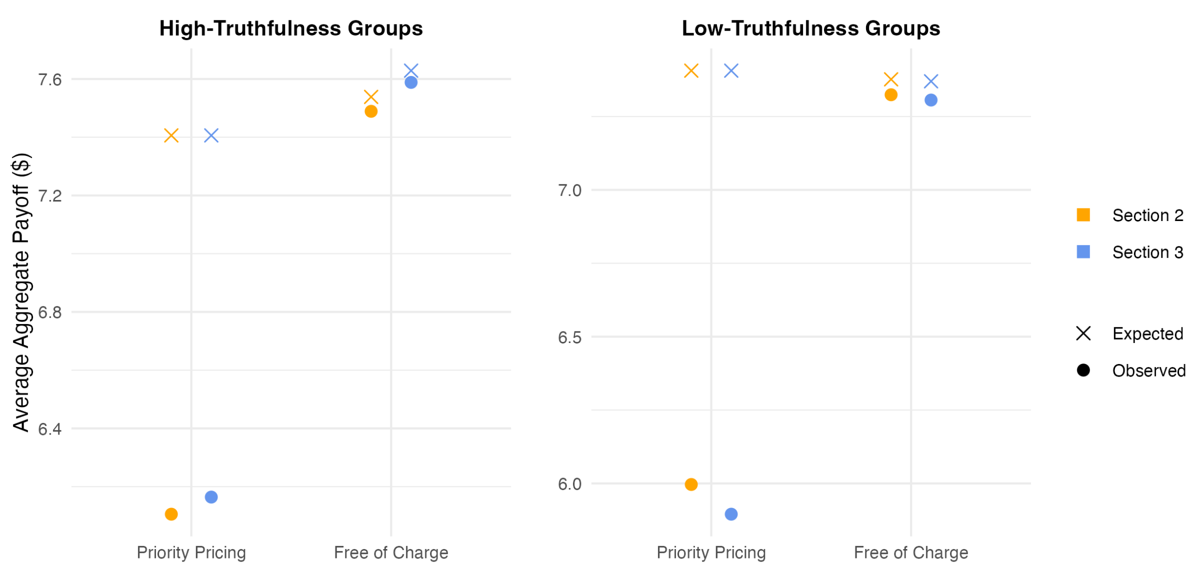

The results presented in Section 5.3 show that beliefs about others' propensity for truthfulness has a significant effect on how participants respond to the truth-telling incentives under priority pricing. Specifically, participants who believed others were more likely to be truthful also tended to be more truthful under priority pricing. Conversely, those who believed others were less likely to be truthful were more inclined to misreport under priority pricing. This behavioral adjustment is significant, as it directly influences the aggregate welfare outcomes.

To illustrate this further, Figure 6 presents a comparison of the average aggregate payoffs for high- and low-truthfulness groups using participant responses from Section 2 and Section 3. This approach allows for a direct comparison of the impact of participant behavior under different informational conditions on aggregate outcomes as the only distinction between Sections 2 and 3 was that, in section 3, participants were explicitly informed about the criteria for qualifying for additional payment. Specifically, in Section 3, they were informed that additional payment criteria required either nine out of ten group members to have been truthful or non-truthful in Section 1, depending on the experimental condition to which they had been assigned.

As shown in Figure 6, in Section 3 the high-truthfulness groups achieved a higher average aggregate payoff compared to Section 2. In contrast, the average aggregate payoff for the low-truthfulness groups in Section 3 was lower than in Section 2. These findings highlight the critical role that beliefs play in shaping individual actions and, consequently, in determining overall aggregate outcomes.

6 Discussion and Conclusion

This study examines the impact of incentive-compatible priority pricing (PP) on consumer welfare and explores how preferences for truth-telling influence these outcomes. The results show that PP generally reduces the incidence of non-truthful reporting among participants. However, contrary to expectations, PP did not have significantly different effects across groups with low versus high underlying propensities for truth-telling. The findings suggest that the lack of significant differences resulted from participants adjusting their reporting behavior based on their perceptions of their group members' truthfulness.

These findings emphasize the importance of considering behavioral factors when designing pricing incentives. Consistent with existing evidence on truth-telling preferences, this study finds that many participants remain truthful even at personal cost. Moreover, not all participants responded to the pricing incentive, resulting in lower average payoffs than theoretical predictions. Given that participants demonstrated understanding of the experimental setup through control questions and those in the priority pricing treatment groups were explicitly informed that truthful responses would maximize their payments, the lack of responsiveness to PP likely stems from factors other than a lack of understanding and is something that warrants further investigation. In fact, this behavioral pattern aligns with the classification by Gneezy et al. (2013), which categorizes individuals into three types: those who never lie, those who always lie, and those who respond to incentives to lie. These findings suggest that future analyses could benefit from incorporating a broader range of truth-telling behaviors.

Moreover, the finding that participants adjust their truth-telling behavior based on their perceptions of others' truthfulness highlights the critical role that beliefs play in the effectiveness of pricing incentives. This observation that participants were more likely to be truthful when they believed others would be, and less so otherwise, aligns with the broader literature on social preferences. This literature suggests that individuals care not only about their own material payoffs but also about social comparisons, fairness, and reciprocity (see, for example, Fehr and Gächter, 2000; Fehr and Schmidt, 2001). However, this study relies on a simplified utility function that does not account for social preferences. Future research could benefit from exploring how incorporating more complex utility functions that account for social preferences might influence the relationship between truth-telling behavior and the effectiveness of pricing incentives.

While this study provides valuable insights, there are two key limitations that should be acknowledged. First, although online data collection offers advantages, such as access to more representative samples and the ability to gather larger sample sizes at a lower cost, it is also more vulnerable to data quality concerns. Research has found that data collected virtually may exhibit higher levels of noise due to participant inattention, which can lead to reduced sensitivity to treatment variations. This issue can be especially pronounced in settings with complex strategic interactions, where repeated interactions and learning are required (see, for example, Fréchette et al., 2022). To mitigate these concerns, in this study, participants were required to successfully complete comprehension checks, and the survey was kept brief, with an average completion time of 12 minutes. Although these measures do not entirely eliminate data quality concerns, it is worth noting that the setup of this experiment was relatively simple, which may help minimize the impact of these issues.

Second, it is important to note that the experimental setup is highly stylized, serving as a proof of concept rather than offering broad generalizability across different contexts. In practice, studies have shown that various nuances influence when and how likely individuals are to be untruthful. Factors such as the magnitude of the lie (Fischbacher & Föllmi-Heusi, 2013) and the size of the financial incentive (Gneezy, 2005) can significantly shape the relevance and magnitude of the effects observed in this study. Despite these limitations, the primary aim here is to identify the general direction of these effects rather than to precisely measure their magnitude.

References

Abeler, J., Becker, A., & Falk, A. (2014). Representative evidence on lying costs. *Journal of Public Economics*, *113*, 96–104. https://doi.org/https://doi.org/10.1016/j.jpubeco.2014.01.005

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

Adiri, I., & Yechiali, U. (1974). Optimal priority-purchasing and pricing decisions in non-monopoly and monopoly queues. *Operations Research*, *22*(5), 1051–1066.

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

Australian Government Department of Health and Aged Care. (n.d.). *About private health insurance*. Retrieved August 15, 2024, from https://www.health.gov.au/topics/private-health-insurance/about-private-health-insurance

Bucciol, A., & Piovesan, M. (2011). Luck or cheating? a field experiment on honesty with children. *Journal of Economic Psychology*, *32*(1), 73–78.

Campos-Mercade, P. (2022). When are groups less moral than individuals? *Games and Economic Behavior*, *134*, 20–36. https://doi.org/https://doi.org/10.1016/j.geb.2022.03.009

Chao, H.-p., & Wilson, R. (1987). Priority service: Pricing, investment, and market organization. *The American Economic Review*, 899–916.

Dolan, R. J. (1978). Incentive mechanisms for priority queuing problems. *The Bell Journal of Economics*, 421–436.

Erat, S., & Gneezy, U. (2012). White lies. *Management science*, *58*(4), 723–733.

Euronews. (2024, July 11). *Germans warned of long passport waiting times: How do other european countries compare?* Retrieved August 15, 2024, from https://www.euronews.com/travel/2024/07/11/germans-warned-of-long-passport-waiting-times-how-do-other-european-countries-compare

Fehr, E., & Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of economic perspectives*, *14*(3), 159–182.

Fehr, E., & Schmidt, K. M. (2001). Theories of fairness and reciprocity-evidence and economic applications. *Available at SSRN 264344*.

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

Fréchette, G. R., Sarnoff, K., & Yariv, L. (2022). Experimental economics: Past and future. *Annual Review of Economics*, *14*, 777–794.

Gerlach, P., Teodorescu, K., & Hertwig, R. (2019). The truth about lies: A meta-analysis on dishonest behavior. *Psychological bulletin*, *145*(1), 1.

Gershkov, A., & Winter, E. (2023). Gainers and losers in priority services. *Journal of Political Economy*, *131*(11), 3103–3155.

Ghanem, S. B. (1975). Computing center optimization by a pricing-priority policy. *IBM Systems Journal*, *14*(3), 272–291.

Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, *95*(1), 384–394.

Gneezy, U., Kajackaite, A., & Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, *108*(2), 419–453.

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

GOV.UK. (n.d.). *Get a passport urgently*. Retrieved August 15, 2024, from https://www.gov.uk/get-a-passport-urgently

GOV.UK. (2024). *Get a faster decision on your visa or settlement application*. Retrieved August 15, 2024, from https://www.gov.uk/faster-decision-visa-settlement

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

Hurkens, S., & Kartik, N. (2009). Would i lie to you? on social preferences and lying aversion. *Experimental Economics*, *12*, 180–192.

Kuchinke, B. A., Sauerland, D., & Wübker, A. (2009). The influence of insurance status on waiting times in german acute care hospitals: An empirical analysis of new data. *International Journal for Equity in Health*, *8*, 1–10.

Lundquist, T., Ellingsen, T., Gribbe, E., & Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior & Organization*, *70*(1-2), 81–92.

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

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

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

OECD. (2020). *Waiting times for health services: Next in line*. https://doi.org/10.1787/242e3c8c-en

Thami, P. (2024, March). Queuing decisions [AEA RCT Registry]. https://doi.org/10.1257/rct.13064-1.0

U.S. Citizenship and Immigration Services. (n.d.). *How Do I Request Premium Processing?* Retrieved August 15, 2024, from https://www.uscis.gov/forms/all-forms/how-do-i-request-premium-processing

U.S. Department of State. (n.d.). *U.S. Passport*. Retrieved August 15, 2024, from https://travel.state.gov/content/travel/en/passports.html

Wilson, R. (1989). Efficient and competitive rationing. *Econometrica*, *57*(1), 1–40. Retrieved May 23, 2024, from http://www.jstor.org/stable/1912571

A Appendix: Incentive Compatible Priority Pricing

This section presents the solution to the designer's optimization problem. The first incentive-compatibility constraint, (IC-1), ensures that it is weakly optimal for the low type agents to truthfully report their type rather than misrepresenting themselves as high type agents. This constraint is satisfied if:

$$\mathbb{E}_{\theta_j}[u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_j)) - p_L] \geq \mathbb{E}_{\theta_j}[u(\theta_L, t(\hat{\theta}_H, \hat{\theta}_j)) - p_H]$$

Equivalently, this can be expressed as:

$$\begin{aligned}&\mathbb{P}[\theta_j = \theta_H](u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_H)) - p_L) + \mathbb{P}[\theta_j = \theta_L](u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_L)) - p_L)\\ &\geq \mathbb{P}[\theta_j = \theta_H](u(\theta_L, t(\hat{\theta}_H, \hat{\theta}_H)) - p_H) + \mathbb{P}[\theta_j = \theta_L](u(\theta_L, t(\hat{\theta}_H, \hat{\theta}_L)) - p_H)\end{aligned}$$

Substituting the respective utilities and probabilities gives:

$$\alpha(\theta_L\delta - p_L) + (1-\alpha)(\frac{1}{2}\theta_L + \frac{1}{2}\theta_L\delta - p_L) \geq \alpha(\frac{1}{2}\theta_L + \frac{1}{2}\theta_L\delta - p_H) + (1-\alpha)(\theta_L - p_H)$$

Simplification yields:

$$p_H - p_L \geq \frac{1}{2}\theta_L(1-\delta) \tag{3}$$

Similarly, the second incentive compatibility constraint, (IC-2), ensures that it is also weakly optimal for high type agents to truthfully report their type rather than misrepresent as low type agents. This constraint is satisfied if:

$$\mathbb{E}_{\theta_j}\ [u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_j)) - p_H] \geq \mathbb{E}_{\theta_j}\ [u(\theta_H, t(\hat{\theta}_L, \hat{\theta}_j)) - p_L]$$

Equivalently, this can be expressed as:

$$\begin{aligned}&\mathbb{P}[\theta_j = \theta_H](u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_H)) - p_H) + \mathbb{P}[\theta_j = \theta_L](u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_L)) - p_H)\\ &\geq \mathbb{P}[\theta_j = \theta_H](u(\theta_H, t(\hat{\theta}_L, \hat{\theta}_H)) - p_L) + \mathbb{P}[\theta_j = \theta_L](u(\theta_H, t(\hat{\theta}_L, \hat{\theta}_L)) - p_L)\end{aligned}$$

Substituting the respective utilities and probabilities, and simplification yields:

$$p_H - p_L \leq \frac{1}{2}\theta_H(1-\delta) \tag{4}$$

The first individual rationality constraint, (IR-1), ensures that the low type agents are not worse off by seeking an appointment. Formally, this condition is given by:

$$\mathbb{E}_{\theta_j}\ [u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_j)) - p_L] \geq 0$$

Expressed equivalently as:

$$\mathbb{P}[\theta_j = \theta_H](u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_H)) - p_L) + \mathbb{P}[\theta_j = \theta_L](u(\theta_L, t(\hat{\theta}_L, \hat{\theta}_L)) - p_L) \geq 0$$

Substituting the respective utilities and probabilities, and simplification yields:

$$p_L \leq \frac{1}{2}\theta_L(1 + \delta - \alpha(1 - \delta)) \tag{5}$$

Similarly, the second individual rationality constraint, (IR-2), ensures that the high type agents are not worse off by seeking an appointment. The condition is expressed formally as:

$$\mathbb{E}_{\theta_j}\left[u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_j)) - p_H\right] \geq 0$$

Equivalently, this can be expressed as:

$$\mathbb{P}[\theta_j = \theta_H](u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_H)) - p_H) + \mathbb{P}[\theta_j = \theta_L](u(\theta_H, t(\hat{\theta}_H, \hat{\theta}_L)) - p_H)$$

Substituting the respective utilities and probabilities, and simplification gives:

$$p_H \leq \frac{1}{2}\theta_H(2 - \alpha(1 - \delta)) \tag{6}$$

Further, the optimal prices need to be non-negative, i.e.:

$$p_L \geq 0 \tag{7}$$

$$p_H \geq 0 \tag{8}$$

Since aggregate utility decreases as p_L and p_H increase, the optimal solution must minimize these payments while satisfying the incentive compatibility (IC), individual rationality (IR), and non-negative prices constraints, namely, equations (3) through (8).

Non-Binding of $p_H \geq 0$

Equation (3), $p_H \geq \frac{1}{2}\theta_L(1 - \delta) + p_L$ (IC constraint for the low type), implies that if $p_H = 0$, then p_L would be negative, violating Equation (7). Therefore, $p_H > 0$.

Binding of $p_L = 0$

When $p_L = 0$ satisfying Equation (7), the IR constraint (Equation 5) for the low type agent is satisfied, and the cost for the low type agents is minimized. For this value of p_L, the smallest p_H that satisfies the IC constraint for the low type (Equation 3) is $p_H = \frac{1}{2}\theta_L(1-\delta)$.

Verifying that $p_H = \frac{1}{2}\theta_L(1 - \delta)$ satisfies all constraints for high type agents

IC constraint (Equation 4) is satisfied as $p_H - p_L = \frac{1}{2}\theta_L(1-\delta)$ and it follows from $\theta_H > \theta_L$ that $\frac{1}{2}\theta_H(1 - \delta) > \frac{1}{2}\theta_L(1 - \delta)$.

IR constraint (Equation 6) is satisfied as $2 - \alpha(1 - \delta) > 1 - \delta$ as $\alpha, \delta \in (0, 1)$ implies $p_H = \frac{1}{2}\theta_L(1 - \delta) < \frac{1}{2}\theta_H(2 - \alpha(1 - \delta))$.

Equation 8 is satisfied as $p_H = \frac{1}{2}\theta_L(1 - \delta) > 0$ as $\theta_L > 0$ and $\delta \in (0, 1)$.

Conclusion

Thus, the optimal pricing scheme is given by setting $p^*_L = 0$ and $p^*_H = \frac{1}{2}\theta_L(1 - \delta)$.

B Appendix: Proofs

The adoption of an incentive-compatible pricing scheme, as opposed to a uniform pricing scheme, affects aggregate expected welfare through two primary mechanisms. First, it affects the allocation of appointment slots by altering the information revelation behavior of θ_{L_n} type agents. Specifically, when a type θ_H agent and a type θ_{L_n} agent simultaneously seek appointments, incentive-compatible pricing ensures that the θ_{L_n} agent truthfully reports their type, allowing the type θ_H agent to consistently receive the first appointment. Without priority pricing incentives, this appointment would be allocated randomly between the two agents.[23] Second, changes in expected aggregate welfare are also driven by differences in the total expected prices paid by agents.

To assess these changes, recall the expression for expected aggregate consumer welfare from equation (1), restated below for convenience:

$$\mathbb{E}_{\theta_i,\theta_j}\left[U(\theta_i,\theta_j)\right] - \mathbb{E}_{\theta_i,\theta_j}\left[P(\theta_i,\theta_j)\right] = \sum_{\theta_i,\theta_j\in\Theta'} \mathbb{P}_{\theta_i\theta_j} U(\theta_i,\theta_j) - \sum_{\theta_i,\theta_j\in\Theta'} \mathbb{P}_{\theta_i\theta_j} P(\theta_i,\theta_j) \tag{1}$$

As discussed above, the impact on expected aggregate welfare arises specifically from changes in the following subset of terms from equation (1):

$$2\mathbb{P}_{\theta_{L_n}\theta_H} U(\theta_{L_n},\theta_H) - \sum_{\theta_i,\theta_j\in\Theta'} \mathbb{P}_{\theta_i\theta_j} P(\theta_i,\theta_j). \tag{4}$$

Under incentive-compatible priority pricing, this expression is given by:

$$2\mathbb{P}_{\theta_{L_n}\theta_H}\ (u(\theta_{L_n}, t(\hat{\theta}_L,\hat{\theta}_H)) + u(\theta_H, t(\hat{\theta}_H,\hat{\theta}_L))) - \sum_{\theta_i,\theta_j\in\Theta'} \mathbb{P}_{\theta_i\theta_j}(p(\hat{\theta}_i) + p(\hat{\theta}_j))$$

where, θ_{L_n} type agents are incentivized to truthfully report their true type as $\hat{\theta}_L$. Substituting the respective utilities and prices, we obtain:

$$\begin{aligned}&2\mathbb{P}_{\theta_{L_n}\theta_H}(\delta\theta_L + \theta_H) - (\mathbb{P}_{\theta_H\theta_H}(p_H^* + p_H^*) + 2\mathbb{P}_{\theta_H\theta_{L_t}}(p_H^* + p_L^*) + 2\mathbb{P}_{\theta_H\theta_{L_n}}(p_H^* + p_L^*)\\ &+ 2\mathbb{P}_{\theta_{L_t}\theta_{L_n}}(p_L^* + p_L^*) + \mathbb{P}_{L_tL_t}(p_L^* + p_L^*) + \mathbb{P}_{\theta_{L_n}\theta_{L_n}}(p_L^* + p_L^*))\end{aligned}$$

Substituting the respective probabilities, setting $p_L^* = 0$ and simplifying yields:

$$2\alpha\gamma(\delta\theta_L + \theta_H) - 2\alpha p_H^* \tag{5}$$

On the other hand, under a uniform pricing scheme with price p, equation (4) is given by:

$$2\alpha\gamma\ (u(\theta_{L_n}, t(\hat{\theta}_H,\hat{\theta}_H)) + u(\theta_H, t(\hat{\theta}_H,\hat{\theta}_H))) - \sum_{\theta_i,\theta_j\in\Theta'} \mathbb{P}_{\theta_i\theta_j}(p(\hat{\theta}_i) + p(\hat{\theta}_j))$$

[23] While the specific allocation of slots under incentive-compatible pricing may also vary when types θ_{L_t} and θ_{L_n}, or when two θ_{L_n} type agents arrive at the time, this does not affect aggregate utility since both agent types have the same valuation (θ_L).

where, θ_{L_n} type agents misrepresent their type as $\hat{\theta}_H$.

It is useful to recall that $\sum_{\theta_i,\theta_j \in \Theta'} \mathbb{P}_{\theta_i\theta_j}(p(\hat{\theta}_i) + p(\hat{\theta}_j)) = \mathbb{E}[p(\hat{\theta}_i) + p(\hat{\theta}_j)]$ and under a uniform pricing scheme $p(\hat{\theta}) = p \ \forall \ \hat{\theta} \in \Theta'$. Therefore, $\mathbb{E}[p(\hat{\theta}_i) + p(\hat{\theta}_j)] = 2p$. Substituting the respective utilities and expected total prices gives:

$$2\alpha\gamma \ (\frac{1}{2}(\theta_L + \delta\theta_H) + \frac{1}{2}(\delta\theta_L + \theta_H)) - 2p. \tag{6}$$

By comparing the expressions in equation (5) and equation (6), the conditions under which the incentive compatible priority pricing scheme will result in an increase in expected welfare outcome compared to uniform pricing schemes can be derived. This analysis leads to the following propositions:

Proposition 1. *When the expected cost for agents under a uniform pricing scheme is less than under the incentive compatible priority pricing scheme, i.e. when the uniform price is set to $p' = \varepsilon\alpha p^*_H$, where $\varepsilon \in (0,1)$, introducing incentive-compatible priority pricing (PP) will increase expected aggregate consumer welfare if and only if the probability of non-truthful, low-type agents (γ) exceeds the threshold $\gamma^\dagger$, where*

$$\gamma^\dagger = \frac{\theta_L(1-\varepsilon)}{\theta_H - \theta_L}.$$

Proof. The expected aggregate utility from an incentive compatible priority pricing scheme will exceed the expected aggregate utility under a uniform pricing scheme with a lower expected cost if (5) is greater than (6) when $p^*_H = \frac{1}{2}\theta_L(1-\delta)$ and $p = \varepsilon\alpha p^*_H$, where $\varepsilon \in (0,1)$, i.e. :

$$2\alpha\gamma(\delta\theta_L + \theta_H) - 2\alpha\frac{1}{2}\theta_L(1-\delta) > 2\alpha\gamma \ (\frac{1}{2}(\theta_L + \delta\theta_H) + \frac{1}{2}(\delta\theta_L + \theta_H)) - 2\epsilon\alpha\frac{1}{2}\theta_L(1-\delta)$$

Simplification gives:

$$\gamma^\dagger > \frac{\theta_L(1-\varepsilon)}{\theta_H - \theta_L}$$

The threshold $\gamma^\dagger$ is well-defined as $\theta_H > 2\theta_L > 0 \implies \theta_H > (2-\epsilon)\theta_L > 0$ as $\varepsilon \in (0,1)$. □

Proposition 2. *When there is no cost to the agents, i.e. when the price is set to $p = 0$, introducing incentive-compatible priority pricing (PP) will increase expected aggregate consumer welfare if and only if the probability of non-truthful, low-type agents (γ) exceeds the threshold γ^*, where*

$$\gamma^* = \frac{\theta_L}{\theta_H - \theta_L}.$$

Proof. The expected aggregate utility from an incentive compatible priority pricing scheme will exceed the expected aggregate utility when there is no cost to the agents if equation (5) is greater than equation (6) when $p_H^* = \frac{1}{2}\theta_L(1-\delta)$ and $p = 0$, i.e. :

$$2\alpha\gamma(\delta\theta_L + \theta_H) - 2\alpha\frac{1}{2}\theta_L(1-\delta) > 2\alpha\gamma\ (\frac{1}{2}(\theta_L + \delta\theta_H) + \frac{1}{2}(\delta\theta_L + \theta_H))$$

Simplification yields:

$$\gamma^* > \frac{\theta_L}{\theta_H - \theta_L}.$$

The threshold γ^* is well-defined as $\theta_H > 2\theta_L > 0$. This result can also be directly derived from Proposition 1, as it corresponds to the case where $\varepsilon = 0$. □

Proposition 3. *When the expected cost for the agents under the incentive-compatible priority pricing (PP) scheme equals that under a uniform pricing scheme, i.e., when the uniform price is set to $p = \alpha p_H^*$, introducing priority pricing (PP) will always generate higher expected aggregate consumer welfare.*

Proof. Substituting the uniform price $p = \alpha p_H^*$ into equation (6) gives:

$$2\alpha\gamma\ (\frac{1}{2}(\theta_L + \delta\theta_H) + \frac{1}{2}(\delta\theta_L + \theta_H)) - 2\alpha p_H^*$$

Equation (5) can be rewritten as:

$$2\alpha\gamma(\frac{1}{2}(\delta\theta_L + \theta_H) + \frac{1}{2}(\delta\theta_L + \theta_H)) - 2\alpha p_H^*$$

The later is greater than the former as $\theta_H > \theta_L$ and $\delta \in (0, 1)$. □