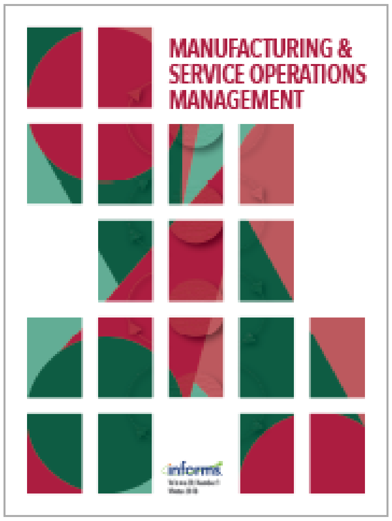

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

No Call, No Show: Impact of No-Shows on Customer Attrition in Online-to-Offline Services

Nil Karacaoglu, Simin Li

To cite this article:
Nil Karacaoglu, Simin Li (2026) No Call, No Show: Impact of No-Shows on Customer Attrition in Online-to-Offline Services. Manufacturing & Service Operations Management

Published online in Articles in Advance 29 Jul 2026

. https://doi.org/10.1287/msom.2025.0421

informs
https://pubsonline.informs.org/journal/msom
MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–20
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

No Call, No Show: Impact of No-Shows on Customer Attrition in Online-to-Offline Services

Nil Karacaoglu,[a] **Simin Li**[b,*]

[a] Fisher College of Business, Ohio State University, Columbus, Ohio 43210; [b] Freeman School of Business, Tulane University, New Orleans, Louisiana 70118

*Corresponding author

karacaoglu.2@osu.edu (N. Karacaoglu), sli41@tulane.edu (S. Li)

https://orcid.org/0000-0003-4883-6870 (N. Karacaoglu), https://orcid.org/0000-0001-8979-502X (S. Li)

Received: June 3, 2025
Revised: March 10, 2026
Accepted: June 2, 2026
Published Online in Articles in Advance: July 29, 2026

https://doi.org/10.1287/msom.2025.0421

Abstract. ***Problem definition*:** Worker absenteeism poses operational challenges across industries. Although extensively studied in traditional workplaces, its impact on online-to-offline service platforms remains overlooked. Unlike team-based settings, where colleagues can cover absences, customers on these platforms are particularly vulnerable to service provider no-shows because each customer is matched with one independent contractor. The service cannot proceed if the provider fails to appear. This paper addresses this gap by examining how service provider no-shows affect customer attrition and investigating the unique challenges and opportunities in designing effective mitigation strategies within the service platform context. ***Methodology/results*:** Using detailed transaction data from a European home-cleaning platform, we find that no-shows increase customer attrition probability by 25.6%, especially among loyal customers, and reduce customer lifetime value by 1.55% despite their low prevalence (3.51%). These notable adverse effects call for strategies to combat no-shows. We find that predicting gig workers' no-shows ex ante is challenging, highlighting the need for post–no-show remedy strategies. Our findings suggest that frustration discounts after no-shows can help reduce attrition, whereas finding a substitute to restore service later the same day does not mitigate attrition. There is suggestive evidence that restoration is effective only if provided promptly (within one hour). We then build an analytical model that takes these empirical insights as inputs and delivers profit-maximizing optimal remedy strategies for service platforms tailored to customer and platform characteristics. ***Managerial implications*:** Our study sheds light on the significant adverse impact caused by no-shows in online-to-offline service settings, despite their rarity. We provide a framework that enables online-to-offline service platforms to design optimal remedy strategies to combat no-shows based on their own data, addressing the unique challenges of worker absenteeism in the gig economy settings.

History: This paper was selected as part of the 1RR initiative between the *M&SOM* Journal and the MSOM Society.
Supplemental Material: The online appendices are available at https://doi.org/10.1287/msom.2025.0421.

Keywords: absenteeism • service failure • online-to-offline platforms • service recovery • customer retention • customer lifetime value

1. Introduction

According to the U.S. Bureau of Labor Statistics, the average absence rate among full-time wage and salary workers was 3.2% in 2021.[1] Although it is widely acknowledged that service provider no-shows degrade customer experience, the extent of the impact can vary significantly across industries. Previous literature has focused on industries where workers collaborate in teams—nurses, retail sales associates, and waiters—where covering for absent colleagues is challenging but possible. In contrast, absenteeism can severely impact service platforms when each customer is matched one-to-one with an independent contractor. With the rising popularity of service platforms, it is not uncommon for customers to bear the brunt of no-shows, such as cleaners not showing up for cleaning.[2]

Although service provider no-shows can be consequential, addressing them is not trivial for platforms. Service providers operate as independent contractors with autonomy over their schedules, leaving platforms with limited control. Additionally, such autonomy might make predicting no-shows more challenging, heightening the need for effective post-failure recovery strategies. Yet, empirical evidence on both the impact of such service failures and the effectiveness of recovery strategies remains scarce in the context of service platforms.

Our study addresses this critical gap by developing a comprehensive framework. We first quantify the effect of no-shows on customer attrition and lifetime value and examine how these effects vary with customer loyalty. We then explore the predictability of no-shows ex ante and estimate the effectiveness of two post-failure recovery strategies. Motivated by our empirical findings, we develop an analytical model to help platforms design optimal remedy strategies tailored to their business context. Our goal is to provide service platforms with actionable insights to combat no-shows.

We conduct our analyses using granular transaction data from a European cleaning platform operating across four countries and eight cities. Similar to U.S.-based Handy, this platform connects individuals seeking cleaning services with cleaners. Our data set includes detailed information on cleaning requests and cleaners' unannounced absences (i.e., no-shows) at the cleaning level, enabling us to analyze how no-shows affect subsequent customer behavior.

The platform's core model, connecting an independent contractor to a customer, mirrors operations in the rapidly growing online-to-home services industry. Valued at USD 3.71 billion globally, this industry is projected to grow at 16.7% annually through 2030.[3] The industry encompasses a wide range of platforms serving diverse needs. For example, customers can request furniture assembly via TaskRabbit, hire electricians and plumbers through Thumbtack, book yardwork through HomeAdvisor, or schedule cleanings with Housekeep. Given the scale and diversity of this industry, our study of service provider no-shows and potential remedies offers timely and practically relevant insights.

We start by showing that no-shows substantially increase customer attrition on online-to-home services platforms. Using logistic regression at the cleaning level, we find that customers are significantly less likely to continue using the platform following a no-show incident, increasing the probability of attrition by 25.6%. To ensure robustness of this finding, we perform additional analyses, including repeating our analysis on a matched sample, controlling for an expanded set of variables, using alternative specifications such as survival models, and running falsification tests. Furthermore, we show that our results cannot be attributed to alternative mechanisms such as disintermediation. We then quantify the financial consequences of no-shows using the customer lifetime value (CLV) framework developed by Bachmann et al. (2021) and find that no-shows reduce customer lifetime value by 1.55%. Despite the relatively low prevalence of no-shows (3.51% of cleanings in our sample), the results indicate that no-shows are costly for the platform.

We also investigate how the impact of no-shows on attrition varies with customer loyalty. Prior literature, which leveraged laboratory experiments and surveys, offers mixed perspectives on whether loyal customers are more or less forgiving of service failures. Using transaction data, we show that in online-to-home services platforms, the effects of no-shows on attrition are more pronounced among loyal customers of the platform. This result carries significant strategic implications, given prior studies highlighting that repeat customers generally represent a substantial source of revenue across industries and that acquiring new customers can cost 5 to 25 times more than retaining existing ones (Gallo 2014). The increased sensitivity of loyal customers to service failures highlights the importance of tailoring no-show mitigation strategies according to customer heterogeneity.

To examine whether the platform can predict no-shows and undertake preventive measures, we employ XGBoost, an advanced machine learning model, and include a comprehensive set of features that capture cleaner, customer, and cleaning characteristics. We find that accurately predicting no-shows is challenging. The limited predictive performance further underscores the importance of developing remedy strategies that can be deployed when no-shows occur. Using the instrumental variable (IV) method, we causally estimate the effectiveness of two remedies offered by the platform: frustration discounts and same-day restoration. Whereas frustration discounts can significantly reduce customer attrition after no-shows, reassigning and completing the service, which we call restoration, later on the same day proves ineffective. We find suggestive evidence that restoration is effective only if performed promptly—within an hour of the originally scheduled time.

Informed by these empirical insights, we develop an optimization model to help platforms design profit-maximizing remedy strategies tailored to customer types. Our model focuses on two empirically motivated remedies: offering discounts following no-shows and retaining standby service providers for prompt restoration. This model delivers a strategy comprising (i) the optimal number of standby service providers and their optimal allocation across customer types and (ii) the optimal frustration discount to offer each customer type when prompt restoration is not feasible.

Collectively, our paper makes the following contributions to academic research and managerial practice.

- Our study examines a gig economy platform where each customer is served by one provider, in contrast to extensively studied team-based settings where coworkers can often cover for absent colleagues. In a one-to-one context, customers directly experience the adverse impacts of no-shows, which allows us to precisely measure the effect of no-shows on customer attrition (Section 4.1) and document that loyal customers react more adversely to service failures (Section 4.3).
- We study gig workers who have autonomy over whether and where to work. This differs from prior research that focused on traditional employees. We show that schedule-related variables are not strong

predictors of no-shows in this setting with worker autonomy (Section 4.2.1), limiting the platform's ex ante preventive measures.

- We causally identify the effectiveness of two remedy strategies for ex post recovery using customer transaction data, contributing to the literature on service failure recovery that predominantly relies on laboratory experiments and surveys measuring purchase intentions (Section 5.1). We further develop an analytical model to design effective post-failure remedies.

- Our study offers practical insights for industry practitioners on two fronts. First, we show that it is important to mitigate the impact of no-shows because, although rare, no-shows are costly to platforms. Second, we develop a comprehensive framework that similar platforms can use. This framework quantifies no-show impacts, estimates cleaning-level no-show probabilities, evaluates remedy effectiveness, and inputs these parameters into an optimization model. This framework helps platforms generate tailored optimal remedy strategies using their own data, as we show in Section 5.2.

- Prior research on service operations predominantly examines process failures (e.g., longer wait times). In contrast, we focus on outcome failures where the promised core service is not fulfilled (scheduled cleaning at a specific date and time). We find that these outcome failures have significant adverse impacts on customers. Additionally, existing research mainly considers short-term outcomes like immediate abandonment, whereas we analyze longer-term effects of no-shows on customer attrition and lifetime value. We elaborate on our contribution to these streams of literature in Section 2.2.

2. Literature Review

Our research contributes to two streams of literature on service management: (i) worker management, absenteeism, and labor staffing; and (ii) service failures.

2.1. Service Provider Absenteeism

Customer no-shows, such as those of patrons and patients, have been widely studied in the literature (Hassin and Mendel 2008, Liu et al. 2019, Kong et al. 2020). However, research on service provider no-shows (i.e., absenteeism) remains limited. Moreover, existing studies on absenteeism primarily focus on settings where providers work in teams, which differs from the gig economy context we examine.

Early research links absenteeism to heavy workloads, stress, and understaffing, especially in healthcare (McVicar 2003, Unruh et al. 2007). Its consequences include lower service quality and reduced productivity across sectors (Unruh et al. 2007, Kesavan et al. 2022). As service providers work in teams in these settings, the strategy to deal with absenteeism has centered around incorporating potential worker absences into ex ante staffing policies. Fry et al. (2006) develop firefighter staffing policies that account for unpredictable absences. Green et al. (2013) propose optimal nurse staffing by accounting for endogenous absenteeism rates. Wang and Gupta (2014) utilize individual absenteeism history in nurse staffing planning.

Recent research has demonstrated the impact of scheduling practices on productivity and absenteeism in traditional employment settings. Studies show that schedule consistency positively affects productivity in retail settings (Kesavan et al. 2022, Lu et al. 2022). Similarly, Kamalahmadi et al. (2021) show that, although short-notice schedules do not harm server productivity overall, real-time schedules reduce productivity significantly. Moreover, Kwon and Raman (2023) provide evidence that inconsistent work schedules significantly increase the probability of lateness and absenteeism. These findings underscore the importance of scheduling and schedule predictability in traditional employment settings where managers have direct control over worker schedules.

Our research addresses a critical gap in the literature by studying the gig economy, where service providers have complete autonomy over when and whether to work. Traditional, ex ante imposed schedules are not viable in this context. As we show in our prediction analysis (Section 4.2.1), schedule-related factors are not strong predictors of no-shows, possibly due to the autonomous nature of gig work. Therefore, instead of focusing on such interventions, we examine two ex post levers that platforms can implement—retaining standby cleaners and offering frustration discounts—that are more applicable to our setting and could help mitigate the adverse impacts of no-shows.

Our paper also contributes to the growing literature on service provider and customer behavior in gig-economy settings. Prior studies have primarily examined gig workers' labor supply decisions (Allon et al. 2023, Xu et al. 2023) and customers' sensitivities to service delays (Yu et al. 2022). We extend this literature by documenting the impact of service provider absenteeism and by developing a framework that enables platforms to effectively mitigate these failures, thereby addressing a crucial operational challenge for online-to-offline gig platforms.

2.2. Service Failure

Our work is also related to streams of literature on service failures and service quality management.

2.2.1. Consequences of Service Failures. Dating back to Zeithaml et al. (1996), the marketing literature has shown that service quality affects individual customer behavior. Service failures notably decrease customers' service evaluations and satisfaction (Anderson et al. 2008, Xu et al. 2019). Although service failures can take multiple forms—including advance cancellations and no-shows—our focus is on service provider no-shows. Such failures are more likely to disrupt customers' plans and expectations compared with advance cancellations.

The psychological mechanisms underlying the relationship between service failures and customer behavior explain how service failures can lead to customer attrition. The expectation disconfirmation theory posits that customers evaluate their service experiences by comparing actual performance to expectations and reference levels (Oliver 1980). When service failures occur, perceived performance falls below customers' expectations, producing negative disconfirmation that translates into negative emotions—such as anger and frustration—and dissatisfaction, heightening the propensity to churn. From a service quality perspective, the widely used SERVQUAL framework demonstrates that reliability is consistently the most critical dimension in determining customers' overall service quality perceptions (Parasuraman et al. 1988). No-shows of service providers directly undermine this reliability dimension, because they represent failures to perform promised services dependably and at the agreed-upon time. This deterioration in perceived service quality can increase customer attrition.

Beyond survey-based measures of emotions and satisfaction, this stream of research has also informed optimal service design models. Hall and Porteus (2000) solve for firms' capacity decisions accounting for customers' responses to service failures. These negative perceptions and dissatisfaction have consequential monetary outcomes. Customers who experience inadequate service time in a personal finance setting or more stockouts in a retail setting spend less with the firm (Oliva and Sterman 2001, Anderson et al. 2006). Similarly, denied and downgraded services on overbooked airlines, longer-than-expected wait times, and slower-than-expected service speeds on ride-hailing platforms reduce customers' future usage and spending (Wangenheim and Bayón 2007, Cohen et al. 2021).

Our research makes several contributions that address key gaps in the service failure literature. Previous literature has classified service failures as process failures (deficiencies in service delivery, such as excessive waiting times) and outcome failures (where the core promised service itself is not delivered) (Sivakumar et al. 2014). The empirical service operations literature has predominantly examined process failures across diverse settings, such as restaurants (Allon et al. 2011). In contrast, our study focuses on *outcome* failures—specifically, cleaning services scheduled for particular dates and times that are not delivered due to provider no-shows.

Our findings reveal that such outcome failures significantly increase customer attrition rates, with more substantial customer impacts than typically expected in process failures. The SERVQUAL framework (Parasuraman et al. 1988) provides theoretical insight into this difference: outcome failures, such as no-shows, breach the reliability dimension of service quality. Process failures, such as delays, primarily erode responsiveness. This distinction explains why outcome failures can have large impacts on attrition, as demonstrated by our findings.

Second, although existing service operations research mainly considers short-term outcomes such as customer abandonment (Yu et al. 2022) and decreases in hourly sales (Mani et al. 2015), we analyze the effects of service failures on longer-term outcomes, specifically attrition and customer lifetime value. Together with Cui et al. (2023), our research is among the few studies that examine longer-term consequences of service failures.

Third, whereas the marketing literature primarily utilizes surveys and laboratory experiments to study the impacts of service failures on satisfaction and purchase intentions, we complement these studies by using customer transaction data, which allows us to observe real purchasing behavior rather than relying on stated intentions and quantify the impact of service failures on subsequent customer attrition causally. This data set also allows us to assess the heterogeneous effects of no-shows across customer loyalty levels more reliably.

2.2.2. Effect of Service Failures on Loyal Customers.

Although service failures generally hurt customer-platform relationships and increase churn, how this effect varies with customer loyalty remains theoretically ambiguous and empirically underexplored.

Competing theoretical perspectives offer contradictory predictions. On the one hand, loyal customers may respond less negatively to service failures and may be less likely to leave the platform in the face of a no-show. Loyal customers with several prior service encounters often carry a reservoir of goodwill. Previous work on service failure and recovery suggests that an ongoing relationship and goodwill can buffer against dissatisfaction because the failure incident is judged against a long run of positive interactions rather than in isolation (Tax et al. 1998, Hess et al. 2003). This buffer effect might operate alongside the halo effect, where customers' overall positive impression of the brand casts a favorable light on specific negative incidents, causing them to interpret failures more generously than they otherwise might (Nisbett and Wilson 1977). These mechanisms collectively suggest that loyal customers should respond less negatively to service failures.

Yet, loyalty can just as plausibly magnify the negative impacts of service failure. Contrast theory suggests that when current performance falls below a high reference point, the gap is perceived as larger; consequently, loyal customers might compare a no-show with their past positive interactions, experiencing a heightened sense of loss (Anderson 1973, Oliver and DeSarbo 1988). Fairness theory reinforces this view: After investing substantial time and money across multiple service encounters, loyal customers develop expectations of fault-free execution, making no-shows feel like violations of psychological contracts that trigger stronger negative emotions (Grégoire and Fisher 2008). In addition, loyal customers may also develop a sense of entitlement (Wetzel et al. 2014, Li et al. 2017), believing the platform owes them reliable service. Thus, a failure may cause resentment, because they might consider a no-show to be a betrayal of trust. This could potentially trigger the "How could this happen to me?" reaction that converts disappointment into exit (Grégoire and Fisher 2008). In light of these mechanisms, loyal customers should respond more negatively to service failures.

Existing empirical evidence remains mixed and is based on limited methodological approaches. Whereas some studies find that loyalty cushions dissatisfaction (Tax et al. 1998, Hess et al. 2003), others report heightened negative reactions among loyal customers (Grégoire and Fisher 2008). However, these findings rely predominantly on scenario-based vignette studies (Hess et al. 2003) and retrospective surveys (Grégoire and Fisher 2008) that capture hypothetical or retrospective evaluations rather than actual customer behavior.

Our study addresses these limitations in several ways. First, our granular data set allows us to assess heterogeneous no-show effects across loyalty levels more precisely than previous studies. Second, our optimal remedy strategy accounts for customer heterogeneity in loyalty levels. Overall, we contribute to the theoretical debate while offering practical guidance for platforms managing loyal customer relationships in the face of service failures.

2.2.3. Mitigation and Remedies. By examining the negative impact of service failures, our work also contributes to the literature that studies service failure recovery. In the marketing literature, monetary compensation and sincere customer scripts (including apologies and explanations) have proven effective in recovering customer satisfaction. Craighead et al. (2004) and Noone and Lee (2011), and the references therein, provide an extensive review of service recovery. Both papers measure the effectiveness of mitigation strategies by surveying customers' perceptions toward the service and their repatronizing intentions postservice recovery activities.

As noted by Cohen et al. (2021), most prior research on service recovery uses surveys and laboratory experiments. Notable exceptions include Anderson et al. (2006) and Cohen et al. (2021), who examine recovery strategies through field experiments. Anderson et al. (2006) find that offering a price discount on the stockout product or telling the customer that the stockout item is extremely popular preserves the order. Cohen et al. (2021) find that offering a $5 ride credit and an apologetic message incentivizes a frustrated customer to return and use the ride-hailing service again.

Our contribution to the service failure mitigation literature is twofold. First, we use transactional data to causally examine the effectiveness of two remedy strategies—monetary discounts and same-day restoration—in reducing customer attrition. Second, we develop an optimization model to determine optimal remedy strategies that maximize aggregate customer lifetime value, a proxy for profit, for service platforms.

3. Empirical Setting and the Data

3.1. Empirical Setting

Our data set comes from a European two-sided cleaning platform that connects households and businesses looking for cleaning services with cleaners. Customers provide detailed information, including their district, date and start time, the number of rooms in their home, and extra services (e.g., oven or fridge cleaning) when they submit the cleaning request. They also provide their credit card information for payment. The platform predetermines the prices for services.

Posted cleanings are shared with all cleaners registered to serve the customer's district. Cleaners are independent contractors—the platform does not control their schedule or cleaning assignments. Once picked up by a cleaner, a service request is no longer visible to other cleaners. The platform handles all communication, and cleaners and customers cannot directly communicate with each other.

On the day of the cleaning, service is set to start at a specified time. If the cleaner does not show up and the platform confirms the no-show, then the platform records this incident as a no-show. The platform does not explicitly impose any monetary penalty on the cleaners due to no-shows and does not restrict cleaners' access to future cleaning requests.

If a no-show occurs, the platform first tries to restore the service to be completed at a later time on the same day with a replacement cleaner. If this does not work, the platform offers to schedule the cleaning for another day. The platform may also offer monetary discounts to compensate for the inconvenience experienced by the customer. If a cleaning is completed, the cleaner is paid after the cleaning.

3.2. Data Description

The platform shared comprehensive data spanning all cleaning requests from its inception in April 2016 to February 2020. Our analyses focus on the cleaning services delivered after the platform started to record absenteeism data in December 2018.

For each cleaning, we observe the order ID; the date and channel (website, call center, or mobile app) through which the cleaning request was posted; the ID and district of the customer who posted the request; the number of rooms in the customer's home; extra services (e.g., oven or fridge cleaning); the date the cleaning occurred; the ID of the cleaner who worked on the cleaning; whether the customer has pets at home; whether the customer used a coupon or discount; the customer's rating and review of the cleaner, if provided; and the cleaner's rating and review of the customer, if provided.

In addition, the platform shared detailed data on no-shows. This data set covers the order ID of the cleaning for which the no-show took place, the ID of the cleaner who was absent,[4] the date of the no-show, and remedial actions taken by the company to alleviate the effects of the no-show, if any.

By combining these two data sets, we can pinpoint the specific cleanings (identified by the cleaning ID) and customers (identified by the customer ID) impacted by no-shows, along with the date of the no-show and any remedial measures implemented by the platform. Observing customers' order history and no-show information at the cleaning level allows us to evaluate the effect of no-shows on customer attrition.

Our analysis sample includes 81,960 cleanings involving 2,664 cleaners and 32,321 customers.[5] In this sample, the no-show rate is 3.51%. We also provide more detailed summary statistics at the customer, cleaner, month-year-city levels in Table 1.

3.3. Variable Definitions

In this study, we focus on exploring the effects of service providers' no-shows on customer attrition. We specifically look at how likely a customer is not to place another order on the platform after experiencing a no-show. Because these no-show events happen at the cleaning level, we base our analysis on the cleaning level. In this section, we define the variables used in our study and provide corresponding summary statistics in Section A.1 in Online Appendix A.

3.3.1. Dependent Variable. We measure customer attrition by examining whether a customer does not place a subsequent service request following the current cleaning. For this purpose, we define the indicator variable R_i, which equals one if the customer does not place another service request after service i during our period of observation and zero otherwise.[6]

3.3.2. Independent Variable. Our independent variable captures the occurrence of a no-show event for a particular cleaning. If the originally matched service provider fails to show up on the scheduled service date, the cleaning is considered to have experienced a no-show. The indicator variable $NoShow_i$ equals one if cleaning i experiences a no-show and zero otherwise.

3.3.3. Controls. To control for cleaning and customer-specific factors that might impact customer attrition, we include a set of controls defined below.

3.3.3.1. Customer Order History. A customer's response to no-shows and their loyalty can be influenced by their previous interactions with the platform, in other words, their order history. For example, customers who received more cleanings from the platform might be more inclined to return to the platform following a cleaning. To account for the potential influence of a customer's order history on their attrition, we incorporate the $CustomerOrderHistory_i$ variable, which represents the total number of cleanings the customer has received from the platform since its inception prior to cleaning i.

3.3.3.2. Price. The cost of the cleaning service may influence a customer's likelihood of coming back to the platform. To capture this potential effect, we introduce the variable $Price_i$, which denotes the total price of cleaning i, encompassing both the primary cleaning charge and any extra services (such as oven or fridge cleaning) before any discounts are applied.

3.3.3.3. Discount Amount. Customers can pay reduced prices by using a coupon or gift card, as well as by giving or receiving referrals. We calculate the discount amount as the total monetary discount applied by the customer to cleaning i before the scheduled service, that is, $DiscountAmount_i$.

Table 1. Summary Statistics for Platform Activity

	Mean	Median	Standard deviation	Q25	Q75	Q90
Total cleanings per customer	2.54	1.00	4.22	1.00	2.00	5.00
Total cleanings per cleaner	30.77	12.00	47.14	4.00	36.00	84.00
Monthly cleanings per customer	1.36	1.00	0.87	1.00	1.00	2.00
Monthly cleanings per cleaner	7.73	6.00	6.99	2.00	11.00	17.00
Monthly cleanings per city	794.36	369.00	1,111.28	136.00	819.25	2,880.20
Monthly customers per city	585.85	269.50	852.67	105.50	501.50	2,241.60
Monthly cleaners per city	103.08	41.00	150.11	23.25	87.75	410.20

Note. As our data set ends on February 13, 2020, we excluded the incomplete February month when calculating monthly-level summary statistics.

4. Evaluating the Impact of No-Shows

In the previous section, we introduced the empirical setting, described the data set, and defined the key variables. In this section, we first outline our empirical methodology and present our main findings. In Section 4.2, we first assess whether no-shows can be predicted ex ante, followed by a series of analyses that confirm the robustness of our main results. Section 4.3 explores the heterogeneous effects of no-shows by examining how their impact varies with customer loyalty. Finally, in Section 4.4, we quantify the financial impact of no-shows on the platform using customer lifetime value analysis (CLV).

4.1. Empirical Strategy and Main Results

To explore the effect of no-shows on customer attrition, we begin by comparing the average customer attrition rates for cleanings with and without no-show incidents. We find that the average attrition rate is 32.7% for cleanings that did not experience no-shows, but the attrition rate jumps to 46.5% when no-shows occur.

Building on this model-free evidence, we estimate a model quantifying how no-shows causally affect customer attrition. Our identification exploits quasi-experimental variation in no-shows arising from the conditional randomness of no-shows given observable characteristics. Section 4.2.1 provides empirical evidence corroborating this identification strategy, as we show that no-shows are difficult to predict, do not follow distinctive patterns, and are not concentrated over a few cleaners or customers. Section 4.2.2 presents comprehensive robustness checks that provide additional evidence that our findings remain consistent under matched samples, alternative sets of controls, model specifications, and alternative mechanisms.

To identify the impact of no-shows on customer attrition, we use the following model:

$$R_i = \beta NoShow_i + \gamma \mathbf{X}_i + \tau_{myc(i)} + \theta_{product(i)} + \epsilon_i, \qquad (1)$$

where i denotes the cleaning ID. The dependent variable R_i is an indicator variable that captures whether customers do not request another cleaning from the platform after cleaning i. Because our dependent variable is binary, we use logistic regression. Our independent variable $NoShow_i$ captures whether a no-show occurred in cleaning i. The control variables are denoted by $\mathbf{X}_i$ and include $CustomerOrderHistory_i$, which enables us to control for customer loyalty, $Price_i$, which controls for the service's price, and the $DiscountAmount_i$, which denotes the total discount amount the customer used when booking cleaning i.

We control for city-specific temporal patterns at the month-year level by including month-year × city fixed effects, $\tau_{myc(i)}$, and account for time-invariant product category-specific characteristics by including product-category fixed effects $\theta_{product(i)}$. We use robust standard errors. Our results are also robust to using clustered standard errors.

We present the results of our analysis in Table 2. All specifications include month-year × city and product-category fixed effects. Column (1) controls for *CustomerOrderHistory*, whereas column (2) additionally includes the *Price* and *DiscountAmount* controls. The coefficient of *NoShow* is positive and statistically significant in both specifications. Based on the estimated effects, experiencing a no-show increases the odds of customer attrition by an average of 34.4%. In other words, when customers experience a no-show, they become more likely not to place another cleaning request through the platform. This estimated coefficient (column (2) of Table 2) further corresponds to a 25.6% increase in attrition probability.[7]

4.2. Prediction and Robustness Analyses

In the previous section, we documented the adverse impact of no-shows on customer retention. In this section, we first examine whether no-shows are predictable and demonstrate that it is difficult to predict no-shows. We then implement a series of analyses to confirm the robustness of our results.

4.2.1. Prediction. In the previous section, we demonstrated that no-shows significantly increase customer attrition. To mitigate this problem, the platform can adopt two strategies: (1) predicting no-shows ex ante to implement preventive measures and (2) applying remedial actions ex post to lessen the adverse effects. Accurately predicting potential no-shows would enable the platform to proactively reduce absenteeism. In this section, we assess whether the platform can accurately predict no-shows, $NoShow_i$, the key independent variable in our main specification.

In our sample, 3.51% of the observations are no-shows. We estimate gradient-boosted trees (XGBoost) because they constitute a state-of-the-art prediction method due to their outstanding performance on sparse data sets.[8]

Table 2. Impact of No-Shows on Customer Attrition

	(1)	(2)
NoShow	0.302***	0.296***
	(0.044)	(0.045)
Observations	81,959	81,959
Log likelihood	−40,432.0	−39,733.1

Notes. All models include month-year × city fixed effects and product-category fixed effects. Column (1) includes *CustomerOrderHistory* control, and column (2) includes *CustomerOrderHistory*, *Price*, and *DiscountAmount* controls. Robust standard errors are in parentheses. Our results are also robust to clustering standard errors at the: (i) date × city, (ii) month-year × city, (iii) customer, (iv) cleaner, (v) two-way clustering: (cleaner and month-year × city), and (vi) two-way clustering: (customer and month-year × city) levels.

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$; $^{+}p < 0.1$.

Our prediction model incorporates a broad set of features that capture various dimensions of cleaning—including characteristics of the cleaner, the customer, and the cleaning service itself—which may impact the likelihood of a no-show. We discuss in detail the prediction features and model selection in Section A.2.5 in Online Appendix A.

The results from our XGBoost model underscore the inherent difficulty of predicting no-shows ex ante. The F1 score of our prediction model is 0.27, much lower than the rule-of-thumb threshold of 0.7, the threshold that indicates strong predictive performance, indicating limited predictive performance. Despite leveraging a comprehensive set of features and a state-of-the-art model for classification tasks, the algorithm struggles to distinguish between cleanings that experienced no-shows and those that did not. As an additional robustness analysis, in Section 4.2.2, we include all the features included in the prediction model in our main equation (Equation (1)), and we find that no-shows increase customer attrition under this specification as well ($\beta = 0.27$, $p < 0.001$). The weak predictive performance of the model suggests fundamental limitations for implementing effective preemptive interventions. This finding highlights the critical importance of developing ex-post remedy strategies to mitigate the impact of no-shows when they occur. We explore the remedy strategies in Section 5.2.

4.2.2. Robustness Analysis. Our empirical strategy leverages quasi-experimental variation in no-shows to identify their causal impact on customer attrition. No-shows appear to occur quasi-randomly, as supported by our prediction analysis, which shows limited predictive performance. To further validate our results, we conduct extensive robustness analyses, summarized in Table 3 with details in Section A.2 in Online Appendix A. Collectively, these analyses confirm that no-shows significantly increase customer attrition.

First, we systematically expand our main specification to ensure that omitted variable bias does not drive our results. Specifically, we estimate our model including (i) cleaner fixed effects to control for time-invariant cleaner characteristics—such as baseline reliability and service quality—that might influence both no-shows and customer attrition; (ii) cleaner fixed effects alongside customer-side controls to account for potential unobserved heterogeneity on both sides of the platform; (iii) customer district fixed effects; (iv) both customer district and cleaner fixed effects; (v) customer district interacted with product category fixed effects; (vi) all prediction features from the prediction model (Section 4.2.1 and Section A.2.5 in Online Appendix A) to capture a wide range of potential confounders; and (vii) date × city fixed effects in place of month-year × city fixed effects to capture city-specific temporal patterns at a more granular level. Our estimates remain stable and statistically significant across all specifications, as reported in Table 3 "Extended set of controls."

Second, because our observation period ends in February 2020, customers who intended to return later

Table 3. Summary of Robustness Checks

Test	Robustness check	Effect size	p-value	Observations
Main model	Baseline specification (main result)	0.30 (0.045)	<0.001	81,959
Extended set of controls	Cleaner FE	0.24 (0.049)	<0.001	81,959
	Cleaner FE + customer-side controls[a]	0.27 (0.053)	<0.001	81,959
	Customer district FE	0.30 (0.046)	<0.001	81,959
	Customer district FE + Cleaner FE	0.25 (0.050)	<0.001	81,959
	Customer district × prod. category FE	0.29 (0.047)	<0.001	81,959
	Prediction Features from Section 4.2.1	0.27 (0.049)	<0.001	81,959
	Date × city FE	0.31 (0.047)	<0.001	81,959
Alternative specifications	Survival analysis (AFT model)	−0.22 (0.05)	<0.001	81,959
	Excluding the last 90 days	0.34 (0.05)	<0.001	64,139
	Alternative attrition definition (90-day inactivity)	0.30 (0.05)	<0.001	64,139
	LPM; Probit	0.08 (0.009); 0.19 (0.03)	<0.001	81,959
Alternative matching specifications	PSM (1:1, caliper = 0.1)	0.27 (0.06)	<0.001	5,022
	PSM (1:3, caliper = 0.1)	0.23 (0.05)	<0.001	9,646
	PSM (1:1, caliper = 0.2)	0.27 (0.06)	<0.001	5,140
Ruling out the alternative mechanism	Controlling for cleaner past no-shows	0.30 (0.05)	<0.001	81,780
Disintermediation	Same-day restoration subsample	0.27 (0.07)	<0.001	80,129
Placebo	Spurious data variation	−0.0008	—	81,959

[a]Customer side controls include (1) customer tenure on the platform (measured in days), (2) total spending prior to the focal cleaning, (3) interactions between customer district fixed effects and product-category fixed effects to account for socioeconomic heterogeneity within districts, (4) an indicator for whether extra services were requested, (5) fixed effects for specific extra service types, (6) an indicator for the presence of pets in the household, (7) start-hour, (8) the number of days between order date and requested service date, and (9) weather conditions on the service day (temperature, precipitation, snow depth, wind speed, and visibility).

might be misclassified as having churned. We address this right-censoring concern through three approaches: (1) survival analysis using accelerated failure time models, (2) excluding cleanings from the final 90 days to ensure all customers have sufficient time to return, and (3) using alternative attrition measures based on 90-day inactivity periods. All approaches yield consistent results, confirming that right-censoring does not explain our findings (Table 3, "Alternative specifications"; Section A.2.2 in Online Appendix A). We further verify the robustness of our results under probit and linear probability models (LPM), and our findings remain unchanged.

Third, we employ propensity score matching to enhance comparability between cleanings that experienced no-shows and those that did not. By matching cleanings on observable characteristics, such as the total number of orders received prior to the focal cleaning (*CustomerOrderHistory*), we create a more balanced comparison that reduces bias from potential confounding factors. We run a series of alternative matching specifications, which are detailed in Section A.2.1 in Online Appendix A. The coefficient for no-shows remains positive and statistically significant under all these specifications (Table 3, "Alternative matching specifications"; Section A.2.1 in Online Appendix A).

Fourth, we run additional analyses to rule out the alternative disintermediation mechanism—where cleaners and customers bypass the platform for direct transactions. A specific concern is that some recorded "no-shows" might represent cases where cleaners arrived but proposed off-platform transactions, which could simultaneously create an apparent no-show and subsequent customer attrition. We address this through two approaches: (1) controlling for individual cleaners' historical no-show patterns, reasoning that systematic disintermediation behavior would be captured in their records and (2) restricting our analysis to no-show incidents followed by same-day restoration through the platform, where disintermediation is implausible because customers who received off-platform service would have little incentive to reschedule through the platform on the same day. Both approaches yield results similar to our main findings (Table 3, "Ruling out the alternative mechanism:Disintermediation"; Section A.2.4 in Online Appendix A). Finally, the placebo tests (see Table 3, "Placebo"; Section A.2.3 in Online Appendix A) confirm that our data structure and model are not spuriously generating observed results.

4.3. Heterogeneous Effect: Customer Loyalty on Customer Attrition

In Section 4.1, we documented the substantial impact of no-shows on customer attrition. As discussed in Section 2.2.2, it is not a priori clear whether the effect of no-shows on customer attrition would be subdued or amplified for the more loyal customers of the platform. In this section, we explore how the impact of no-shows varies with customer loyalty using the following model:

$$\begin{aligned} R_i = {}& \beta_1 NoShow_i + \beta_2 NoShow_i \times CustomerOrderHistory_i \\ & + \beta_3 CustomerOrderHistory_i + \gamma \mathbf{X}_i + \tau_{cmy(i)} \\ & + \theta_{product(i)} + \epsilon_i. \end{aligned} \tag{2}$$

Similar to our main model described in Section 4.1, this model includes month-year × city fixed effects and product-category fixed effects and is estimated using logistic regression. The interaction term $NoShow_i \times CustomerOrderHistory_i$ captures how the impact of no-shows varies with customer loyalty. The result of this analysis is given in Table 4. Column (2) of Table 4 includes *Price* and *DiscountAmount* controls.

The negative and significant coefficient on *CustomerOrderHistory* indicates that, holding no-show status constant, each additional prior cleaning is associated with a reduction in the odds of attrition. Thus, customers with more completed cleanings exhibit lower baseline attrition risk.

We are interested in whether the effect of a no-show on attrition varies with customers' prior order history. The positive and significant coefficient on *NoShow* × *CustomerOrderHistory* indicates that the impact of a no-show on attrition increases with the number of prior cleanings. Together with the positive coefficient on *NoShow*, this implies that, although a no-show raises the odds of attrition at any given level of prior cleanings, the proportional increase in those odds becomes larger as the total number of prior cleanings increases. In other words, whereas customers with longer order histories have lower baseline attrition, the attrition penalty triggered by a no-show is stronger for them.

This finding is particularly important in light of prior research that underscores the strategic value of retaining loyal customers, because they are the key drivers of long-term profitability (Grégoire and Fisher 2008).

These results highlight the adverse effect of no-shows but do not directly measure the financial implications

Table 4. Heterogeneous Effects of No-Shows on Customer Attrition

	(1)	(2)
NoShow	0.16**	0.16**
	(0.06)	(0.06)
CustomerOrderHistory	−0.29***	−0.26***
	(0.01)	(0.01)
NoShow × *CustomerOrderHistory*	0.09***	0.08***
	(0.02)	(0.02)
Observations	81,959	81,959
Log likelihood	−40,415.5	−39,718.8

Notes. All models include month-year × city and product-category fixed effects. Column (2) also includes *rice* and *DiscountAmount* controls. Robust standard errors are in parentheses.
***$p < 0.001$; **$p < 0.01$; *$p < 0.05$; $^{+}p < 0.1$.

for the platform. To quantify the monetary impact of no-shows, we turn to a CLV model in the next section, which allows us to estimate the revenue losses associated with these service failures by examining each customer's purchase history.

4.4. Financial Impact: CLV

Quantifying the financial impact on the platform is crucial for performing a cost-benefit analysis for potential mitigation strategies discussed in detail in Section 5.2.1. To this end, we leverage a widely used CLV model that (1) extends beyond attrition probability by jointly modeling purchase and attrition behaviors, which could be influenced by no-show incidents, and (2) estimates mean spending per transaction. Using this model, we conduct a counterfactual analysis to evaluate the monetary value the platform could have retained had no-shows *not* occurred.

4.4.1. CLV Model. To model customers' repeat purchases and attrition, we use an extension of the Pareto/negative binomial distribution model (PNBD) developed by Bachmann et al. (2021). PNBD is the most widely used framework among researchers and practitioners in nonsubscription settings like ours (Schmittlein et al. 1987, Oblander and McCarthy 2021). Customers can be in one of two states: "Alive" or "Dead" (churned), and the framework jointly models these two processes. When alive, customer i's purchases (cleaning requests in our case) follow a nonhomogeneous Poisson process with time-varying rate $\lambda_i(t) = \lambda_0 \exp(\gamma_p' \mathbf{X}_{it}^P)$, and their lifetime (until churn) follows an exponential distribution with rate $\mu_i(t) = \mu_0 \exp(\gamma_a' \mathbf{X}_{it}^A)$. The baseline purchasing and attrition rates are independently distributed according to gamma distributions: $\lambda_0 \sim \Gamma(r, \alpha)$, and $\mu_0 \sim \Gamma(s, \beta)$. The parameters γ_p, γ_a are coefficients of the covariate vectors $\mathbf{X}_{it}^P$ and $\mathbf{X}_{it}^A$, respectively. We include whether customer i experienced a no-show by time t, $HadNoShow_{it}$, in both $\mathbf{X}_{it}^P$ and $\mathbf{X}_{it}^A$ to investigate the impact of no-shows. The estimated parameter set $\mathbf{\Theta} := (r, \alpha, s, \beta, \gamma_p, \gamma_a)$, together with covariate values $\mathbf{X}_{it}^P$ and $\mathbf{X}_{it}^A$, determines each customer's *expected number of purchases* during a given period. Notably, no-show incidents, among other covariates, affect both purchasing and attrition via the estimated coefficients $\gamma_{p,HadNoShow}$, $\gamma_{a,HadNoShow}$.

We estimate mean spending per transaction through the Gamma-Gamma model proposed by Fader et al. (2005) and Colombo and Jiang (1999). This model allows us to compute the expected mean spending per transaction conditional on individuals' *observed* number of transactions and average spending per transaction. We assume no-shows do *not* affect customers' spending per transaction, as it is primarily determined by the characteristics of customers' homes in our setting. Each customer's lifetime value is then calculated as the product of their expected number of purchases and expected mean spending per transaction. We estimate the CLV model using data from weeks 1 to 52 and predict customer lifetime value over the subsequent eight-week *holdout* period, which extends to the end of the data set. Comprehensive details on the CLV model's implementation, estimation procedures, goodness-of-fit checks, and estimation results are provided in Online Appendix B.

4.4.2. Estimation Results and Counterfactual. Our primary parameters of interest are the impact of experiencing a no-show on customers' lifetime duration, purchasing frequency, and ultimately their lifetime value. Recall that customers' cleaning requests follow a Poisson process. The negative coefficient on $\gamma_{p,HadNoShow}$ (−0.109; see Table B.1 in Online Appendix B) indicates longer interpurchase intervals, leading to fewer cleaning requests in a given lifetime. Customers' lifetimes are exponentially distributed. The positive coefficient on $\gamma_{a,HadNoShow}$ (0.229; see Table B.1 in Online Appendix B) implies that no-shows accelerate the attrition rate, resulting in shorter customer lifetimes. Taken together, we expect customers' lifetime value to decline following a no-show incident.

We then quantify the reduction in CLV due to no-shows through a counterfactual analysis. Specifically, we assume that no-shows are fully eliminated over a six-month period starting in week 28, during which no-shows otherwise constitute 3.85% of cleanings. We compare predicted customer value during the eight-week holdout period under two scenarios: (1) with no-shows and (2) with no-shows fully eliminated. We find that eliminating no-shows increases CLV by 1.55%. On a per-customer basis, each no-show erodes CLV by approximately 40.2% (≈1.55%/3.85%). This result underscores the need for effective no-show mitigation strategies.

5. Mitigating the Impact of No-Shows

5.1. Effectiveness of No-Show Remedy Strategies

In the previous section, we showed that no-shows significantly amplify customer attrition, and this effect is even stronger for more loyal customers on the platform. To mitigate the negative impact of no-shows, the platform implements two remedy strategies: (i) issuing a frustration discount and (ii) same-day restoration. In this section, we discuss these two remedies in detail and evaluate their effectiveness.

5.1.1. Frustration Discount. In addition to coupons, gift cards, and referral discounts, the platform may, at times, provide a discount to the customer following a negative service experience, such as a no-show, low-quality cleaning, or damage to an item in the house. This discount is referred to as a frustration discount, which is captured in the *FrustrationDiscountAmount* variable. The mean value of this variable is 3.71 local currency (standard deviation (SD) = 21.5). Conditional on receiving a frustration discount, the mean of *FrustrationDiscountAmount* is 79.7 local currency (SD = 62.6). In our sample, 20.8% of no-show cleanings receive a frustration discount.[9]

5.1.2. Same-Day Restoration. To alleviate the inconvenience customers experience due to a no-show, the platform may attempt to arrange a cleaner who would complete the cleaning later on the same day as the original cleaning. We refer to this as a same-day restoration. The same-day restored cleaning variable, *SameDayRestored*, is an indicator variable that takes a value of one if a no-show cleaning is completed on the originally scheduled date and zero otherwise. In our sample, 36.3% of the no-show cleanings are eventually completed the same day.

In this section, we evaluate the effectiveness of these two remedy strategies in mitigating the effects of no-shows on customer attrition. To analyze whether implementing frustration discounts or same-day restoration is effective, we interact each of the remedies, *FrustrationDiscountAmount*$_i$ and *SameDayRestored*$_i$, with the *NoShow*$_i$ variable. We include the same set of controls as in Equation (1). The results of these analyses are presented in columns (1) and (2) of Table 5.

Directly estimating the effect of these remedy strategies can lead to biased results. For example, if the platform is more likely to offer these remedies to customers for whom it believes they will be more effective, the effectiveness of the remedies might be overestimated. Conversely, if the platform offers these remedies when it believes customers are more likely to leave, the effectiveness might be underestimated. To address this potential endogeneity bias, we employ the instrumental variable (IV) method.

5.1.3. Endogeneity of Frustration Discount. Following previous literature (Allon et al. 2023, Xu et al. 2023), we employ an instrumental variable approach using Hausman-type IVs (Hausman 1996) to address the endogeneity of the frustration discount variable. We use the *FDOtherCityDate* variable as an instrument for *FrustrationDiscountAmount*.

This instrumental variable captures the average frustration discount issued to customers in all cities within the same country, on the same day of the week, and within the same month, excluding the focal date–city.[10]

For this IV to be valid, it must satisfy the exclusion restriction and relevance conditions. In our setting, the complaints arising from negative service experiences, such as no-shows, are handled centrally by the platform's country-level customer support team, and the frustration discounts are issued by this support team to customers. The exclusion restriction condition requires that our IV, *FDOtherCityDate*, will impact our outcome variable only through its impact on the *FrustrationDiscountAmount*, and this is plausible for several reasons.

First, the customer support team operating on a given day of the week within a month-year is determined in advance and is independent of service outcomes in any particular city on a given day, making it independent of no-shows and customer behavior. Second, we further construct the instrument in a leave-one-out fashion by excluding the focal date-city. As a result, variation in the support team's aggregate discounting tendency is orthogonal to local attrition shocks on the focal date–city. Hence, the instrument should affect customer attrition only through variation in frustration discounts.

Table 5. Effectiveness of Remedy Strategies

	Without IV		With IV	
	(1)	(2)	(3)	(4)
NoShow	0.275***	0.309***	0.466**	0.283*
	(0.049)	(0.056)	(0.176)	(0.132)
FrustrationDiscountAmount	0.005***		−0.048**	
	(0.0005)		(0.017)	
NoShowFrustrationDiscount	−0.003*		0.029†	
	(0.001)		(0.016)	
SameDayRestored		−0.028		0.043
		(0.089)		(0.334)
Observations	81,959	81,959	81,959	81,959
Log likelihood	−39,676.3	−39,740.2	−39,661.5	−29,642.9

Notes. All models include *CustomerOrderHistory*, *Price*, and *DiscountAmount* controls, as well as month-year × city fixed effects and product-category fixed effects. Robust standard errors are in parentheses. For our IV estimates, we also implement a cluster bootstrap at the date × city level with 500 replications, and our results remain robust to inference based on these bootstrap standard errors. Our IV results are also robust to clustering standard errors at the (i) date × city and (ii) month-year × city levels.

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$; †$p < 0.1$.

We further assess the robustness of our findings to possible violations of the exclusion restriction using two additional analyses. First, we include additional controls to capture city-wide shocks that may correlate with the instrument, thereby strengthening the plausibility of instrument exogeneity. Second, we implement Conley bounds, which allow the instrument to have a direct effect on the outcome and evaluate how sensitive our conclusions are to deviations from strict exogeneity. Our findings are robust under both analyses. Details are provided in Section C.2 in Online Appendix C.

Our IV also satisfies the relevance condition because both the endogenous variable *FrustrationDiscount Amount* and the instrumental variable *FDOtherCityDate* are influenced by the support team's frustration discount issuing behavior, including the magnitude of discounts issued. The first-stage regression confirms this relevance, with a statistically significant coefficient for the *FDOtherCityDate* IV ($\beta_{f1,2} = 0.251$, $p < 0.001$ in Equation (3a)).

To build the IV for the interaction term (*NoShow* × *FrustrationDiscountAmount*), we interact our instrument *FDOtherCityDate* with the *NoShow* variable following Wooldridge (2010, p. 133). The first-stage coefficient for the *NoShow* × *FDOtherCityDate* IV is also significant ($\beta_{f2,2} = 2.196$, $p < 0.001$ in Equation (3b)). We implement the Sanderson-Windmeijer (SW) *F*-tests to further assess the relevance of the instruments. The SW *F*-statistics for *FDOtherCityDate* and *NoShow* × *FDOtherCityDate* are 37.4 and 130.1, respectively. The first-stage regression results are presented in Table C.1 in Section C.1, Online Appendix C.

Using these IVs, we estimate the two-stage residual inclusion (2SRI) model. The 2SRI method is preferred over the traditional 2SLS model when the second-stage estimation is nonlinear, as it produces consistent estimates (Terza et al. 2008, Roy et al. 2021, Pendem et al. 2022). The first-stage regressions are in Equations (3a) and (3b), and the second stage regression is in Equation (4). The mean of our instrument *FDOtherCityDate* is 3.73 local currency (SD = 2.87).

$$\begin{aligned} FrustrationDiscountAmount_i = {} & \beta_{f1,1} NoShow_i \\ & + \beta_{f1,2} FDOtherCityDate_i \\ & + \beta_{f1,3} NoShow_i \\ & \times FDOtherCityDate_i \\ & + \gamma_{f1}\mathbf{X}_i + \varepsilon_{1,i}, \end{aligned} \tag{3a}$$

$$\begin{aligned} NoShowFrustrationDiscount_i = {} & \beta_{f2,1} NoShow_i \\ & + \beta_{f2,2} NoShow_i \\ & \times FDOtherCityDate_i \\ & + \beta_{f2,3} FDOtherCityDate_i \\ & + \gamma_{f2}\mathbf{X}_i + \varepsilon_{2,i}, \end{aligned} \tag{3b}$$

where $NoShowFrustrationDiscount_i = NoShow_i \times Frustration DiscountAmount_i$.

$$\begin{aligned} R_i = {} & \beta_{s,1} NoShow_i + \beta_{s,2} NoShow_i \\ & \times FrustrationDiscountAmount_i \\ & + \beta_{s,3} FrustrationDiscountAmount_i \\ & + \beta_{s,4} NoShow\_Z\_resid_i + \beta_{s,5} Z\_resid_i + \gamma_s \mathbf{X}_i + \epsilon_i, \end{aligned} \tag{4}$$

where *Z_resid* denotes the residuals of *FrustrationDiscountAmount* from Equation (3a). They are computed as $Z\_resid_i = FrustrationDiscountAmount_i - \widehat{FrustrationDiscount Amount}_i$, where $\widehat{FrustrationDiscountAmount}_i$ is the predicted value of $FrustrationDiscountAmount_i$ from Equation (3a). Similarly, *NoShow_Z_resid* denotes the residuals of *NoShowFrustrationDiscount* from Equation (3b). They are computed as $NoShow\_Z\_resid_i = NoShow FrustrationDiscount_i - \widehat{NoShowFrustrationDiscount}_i$, where $\widehat{NoShowFrustrationDiscount}_i$ is the predicted value of $NoShowFrustrationDiscount_i$ from Equation (3b). $\mathbf{X}$ includes control variables *CustomerOrderHistory*, *Price*, *DiscountAmount*, month-year × city fixed effects, and product-category fixed effects, as in our main model. Because *FrustrationDiscountAmount* and *NoShowFrustrationDiscount* are continuous, the first-stage models are estimated by ordinary least squares. Consistent with the main specification, the second stage is estimated using a logistic regression.

In the model without the IVs, the coefficient for *FrustrationDiscountAmount* is positive and significant ($\beta = 0.005$, $p < 0.001$). However, directly interpreting this coefficient risks endogeneity bias and could lead to the incorrect conclusion that frustration discounts increase customer attrition. The 2SRI estimation corrects this bias—the coefficient for the *FrustrationDiscountAmount* ($\beta_{s,3} = -0.048$, $p = 0.005$) is negative and significant as shown in column (3) of Table 5. Furthermore, the summation of the coefficients for *FrustrationDiscountAmount* and *NoShow* × *FrustrationDiscountAmount* is negative and significant ($\beta_{s,3} + \beta_{s,2} = -0.020$, $p = 0.052$). These results suggest that in the face of no-shows, providing customers with frustration discounts can help mitigate the risk of customer attrition.

5.1.4. Endogeneity of Same-Day Restoration. We further address potential endogeneity in the same-day restoration variable. Similar to Allon et al. (2023), we utilize coworker information as our instrument. Specifically, we define *coworkers* as all workers operating in the same city, on the same day, starting a cleaning at the same hour. We use the variable *OccupancyOtherDistricts* as an instrument for *SameDayRestored*. The *OccupancyOtherDistricts* variable measures the number of coworkers who are occupied in districts other than the focal cleaning's district within the same city.

OccupancyOtherDistricts plausibly satisfies the exclusion restriction for several reasons. First, cleanings are scheduled at least one day in advance and prices are fixed, so demand is known beforehand, and cleaners independently choose when to work and which cleaning requests to accept. As a result, higher-demand days are unlikely to induce a systematically different composition of cleaners or alter service quality in ways that directly affect customer attrition. Second, *OccupancyOtherDistricts* is constructed in a leave-one-out fashion, so idiosyncratic demand or quality shocks in the focal district—factors that could directly influence the focal customer's attrition decision—do not impact the instrument. Third, we conduct a placebo test to examine whether *OccupancyOtherDistricts* predicts customer attrition in the absence of no-shows. If the instrument violated the exclusion restriction by capturing broader platform conditions that directly affect attrition, it would also predict attrition for cleanings without no-shows, where no restoration decision is involved. Supporting the exclusion restriction condition, we show that in this placebo sample of cleanings that are not impacted by a no-show, *OccupancyOtherDistricts* does not predict customer attrition ($\beta = 0.00059$, $p = 0.35$).

OccupancyOtherDistricts also satisfies the relevance condition: when more coworkers are occupied in other districts, less local slack is available for arranging a same-day restoration following a no-show. The first-stage regression confirms this relevance, with a statistically significant coefficient for the *OccupancyOtherDistricts* IV ($\beta_{fr,2} = -0.013$, $p < 0.001$ in Equation (5)). The mean of our instrument *OccupancyOtherDistricts* is 17.4 (SD = 18.9).

Using the *OccupancyOtherDistricts* IV, we estimate a 2SRI model as follows.[11] The first-stage and second-stage equations are given in Equations (5) and (6), respectively:

$$\begin{aligned} SameDayRestored_i = \beta_{fr,1} NoShow_i \\ + \beta_{fr,2} OccupancyOtherDistricts_i \\ + \gamma_{fr} \mathbf{X}_i + \varepsilon_{r,i}, \end{aligned} \tag{5}$$

$$\begin{aligned} R_i = \beta_{sr,1} NoShow_i + \beta_{sr,2} SameDayRestored_i \\ + \beta_{sr,3} W\_resid_i + \gamma_{sr} \mathbf{X}_i + \epsilon_{r,i}, \end{aligned} \tag{6}$$

where *W_resid* denotes the residuals for the *SameDayRestored* from Equation (5). They are computed as $W\_resid_i = SameDayRestored_i - \widehat{p}_i$, where $\widehat{p}_i$ is the predicted probability of $SameDayRestored_i$ obtained from Equation (5). Similar to our main specification, the control variables are *CustomerOrderHistory*, *Price*, *DiscountAmount*, month-year × city fixed effects, and product-category fixed effects. Because both *SameDayRestored* and R are binary, following Terza et al. (2008), we implement the 2SRI method, estimating logistic regressions in both stages.

To evaluate the relevance of the IV, *OccupancyOtherDistricts*, we implement the Sanderson-Windmeijer (SW) F-test. Although this test is designed for continuous endogenous variables, we follow previous studies (Kamalahmadi et al. 2021) by treating our binary endogenous variable as continuous. The first-stage F-statistic is 41.87 ($p < 0.001$), which exceeds the conventional threshold of 10, indicating that the instrument is not weak. Table C.1 in Section C.1 in Online Appendix C presents the first-stage regression results.

In the model without IVs, the coefficient for *SameDayRestored* is not significant ($\beta = -0.028$, $p = 0.75$). This result is consistent with the 2SRI estimation, where the coefficient for *SameDayRestored* remains insignificant ($\beta_{sr,2} = 0.043$, $p = 0.89$) as shown in column (4) of Table 5. These findings suggest that, in the face of no-shows, simply finding a substitute cleaner on the same day does not effectively mitigate customer attrition risk.

5.1.5. Prompt Restoration.

Our previous result suggests that same-day restoration does not appear to mitigate customer attrition following no-shows. This finding suggests that restorations that occur later in the day may still disrupt customers' plans and hence may not serve as effective remedies. However, prompt restoration—one where the platform urgently secures a substitute cleaner to preserve the customer's original appointment window—can potentially be effective. To assess whether prompt restoration is effective, we explore whether reassigned cleanings (performed by a substitute cleaner) that start within one hour of the original scheduled time can mitigate customer attrition. However, securing a substitute cleaner in such a short time frame is challenging. In our data set, the platform finds a substitute within one hour in only 6% of no-show incidents. To analyze the impact of prompt restoration on customer attrition following a no-show, we replace *SameDayRestored* with *OneHourRestored* in Equations (5) and (6). The binary variable $OneHourRestored_i$ equals one if the reassigned cleaning began within one hour of the original scheduled time and zero otherwise. We follow the same IV strategy discussed above. Our results indicate that the coefficient for *OneHourRestored* is negative and marginally significant in the second stage ($\beta = -1.89$, $p = 0.09$). These results suggest that the effectiveness of restoration may depend on timing: Prompt substitutions can preserve the original appointment window, whereas later restorations may still cause disruption.

5.2. Optimal Remedy Strategy

The results presented in Section 5.1 underscore the potential of remedy strategies to recover customers impacted by no-shows. Specifically, column (3) of Table 5 suggests that offering frustration discounts to customers who experience a no-show can help mitigate subsequent attrition. Given the relative ease of offering discounts, this may be a practical strategy for platforms seeking to retain customers affected by no-shows.

Service restoration presents a more complex picture. Our analysis suggests that same-day restoration, by itself, does not appear to mitigate attrition. Instead, the benefits of restoration seem to depend on timing: prompt restoration may help reduce the negative impact of a no-show. Because the platform relies on gig workers without fixed schedules, prompt restoration is rare, occurring for only 6% of no-shows in our data.

To better achieve prompt substitutions, we propose implementing a standby worker pool, from which workers can be approached ahead of time for potential last-minute assignments in case a no-show occurs on a specific day. Although cleaners' availability on the day remains uncertain (as they are independent contractors), preidentifying such workers increases the likelihood of prompt substitution. To incentivize participation, the platform can compensate standby workers regardless of whether they are eventually called upon. Ultimately, worker availability hinges on the compensation offered for reserving their time.

We develop an optimization model to help platforms manage no-shows through the above-mentioned strategies: offering frustration discounts and managing standby cleaners. Because cleaners typically handle only one job per day, the natural planning horizon is at the *day* level. The sequence of events unfolds as follows. Cleanings requested on a focal day are revealed and picked up by cleaners. The platform then forms a belief about the distribution of no-shows for that day and determines the optimal strategy accordingly. The platform seeks to maximize the aggregate customer lifetime value of all customers to be served on that day, a proxy for profit. They solve the maximization problem to determine the optimal strategy—asking a certain number of cleaners to stand by, offering frustration discounts, or combining both. The uncertainty in daily cleaner absenteeism introduces stochasticity, making the determination of the optimal strategy nontrivial.

Our empirical findings motivate several key features of the model. We introduce the notation here and provide formal definitions in Section 5.2.1. First, customers are more likely to leave the platform after experiencing a no-show, as shown in Section 4.1. This increase in attrition lowers customer lifetime value, as shown in Section 4.4, which we capture in the model through the term $v_{\tau_c}(1-\beta_{\tau_c})$. Second, our results suggest that the effect of no-shows is heterogeneous across loyalty levels: The attrition penalty associated with a no-show is stronger for loyal customers (Section 4.3). We capture this heterogeneity by classifying customers into high- and low-loyalty types, $\tau_c \in \{H, L\}$. Third, frustration discounts may help mitigate the impact of no-shows. Specifically, issuing a nonzero discount d reduces the post–no-show attrition probability (Section 5.1), which we model as $v_{\tau_c}(1-\beta_{\tau_c}(d;\psi_{\tau_c}))$, where $\beta'_{\tau_c}(d;\psi_{\tau_c}) < 0$. Finally, as prompt restoration may help reduce attrition after a no-show (Section 5.1), we assume that a no-show cleaning recovered promptly by a standby cleaner restores the affected customer.

5.2.1. Model Framework. Consider L cleanings to be served on a day. A fraction $\alpha \in (0,1)$ end up as no-shows, where $\alpha \sim F_\alpha$, with mean μ_α and standard deviation σ_α. In practice, a platform can estimate these moments based on our proposed prediction model.[12] We incorporate customer heterogeneity by allowing lifetime values to take one of two types: v_H for high (lifetime) value customers[13] and v_L for low-value customers, where $v_H > v_L$. Let $\tau_c \in \{H, L\}$ denote customer type. A fraction γ of cleanings are requested by high-value customers, and the rest, $1-\gamma$, by low-value customers.

For the $(1-\alpha)L$ cleanings whose cleaner shows up, the platform retains the full customer (lifetime) value, v_{τ_c}, and earns a per-cleaning margin of $p-m$, where p is the price paid by the customer and m is the commission paid to the cleaner. We assume p and m are homogeneous across cleanings. Among the αL no-shows, $\gamma\alpha L$ involve high-value customers, and the remaining involve low-value customers. See the "Value from cleanings not affected by no-show" term in Model (7).

Suppose the platform pays a unit cost μ per cleaner to retain access to standby cleaners. Because cleaners are independent contractors, their actual availability remains uncertain on the focal day. Let β_s denote the probability that a standby cleaner is available when called upon. If the platform retains $\overline{C}$ standby cleaners, the number who are actually available to substitute for absent cleaners follows a binomial distribution: $C \sim Bin(\overline{C}, \beta_s)$, which is realized on the focal day.

When a no-show occurs, the platform can assign a standby cleaner promptly if one is available. In that case, the cleaning proceeds, the platform collects a price p, and the customer's lifetime value v_{τ_c} remains intact. The platform pays an additional commission m_o on top of m, to the standby cleaner. The available standby cleaners are shared among all customers who experienced no-shows. Therefore, the allocation of available standby cleaners to high-value customers cannot exceed $\min\{\gamma\alpha L, C\}$. The allocation to low-value customers, $z \geq 0$, is limited by both their number of no-show incidents, $(1-\gamma)\alpha L$, and the remaining standby cleaners, $C-y$, where $y \geq 0$ denotes the allocation to high-value customers. See the two k_{o,τ_c} "Type-H, standby" and "Type-L, standby" terms in Model (7). Note that our model does not assume high-value customers always receive priority in standby cleaner assignment; the allocation is endogenously determined by the parameters. If no standby cleaner is available, no-show materializes; the focal cleaning does not take place on the originally scheduled date. The platform then decides whether and how much frustration discount to offer to the affected customer.

If the platform does not offer a discount following a no-show, the customer leaves the platform with probability $\beta_{\tau_c} < 1$. We allow for heterogeneous attrition rates across high- and low-value customers in response to no-shows. In this case, the future customer lifetime value becomes $v_{\tau_c}(1-\beta_{\tau_c})$.[14] If the platform offers a frustration discount, $g(d)$, where d denotes the discount amount, full customer satisfaction is still not guaranteed. See the "Type-H, discounts" and "Type-L, discounts" terms in Model (7). We model this by allowing attrition risk to decrease to $\beta_{\tau_c}(d;\psi_{\tau_c}) \in [0, \beta_{\tau_c}]$, where ψ_{τ_c} captures heterogeneity in customers' responsiveness to frustration discounts. We assume $\beta_{\tau_c}(d;\psi_{\tau_c})$ decreases in ψ_{τ_c} for all $d > 0$. As a result, the optimal discount $d^*_{\tau_c}$ is also type dependent. The lifetime value of a customer becomes $v_{\tau_c}(1 - \beta_{\tau_c}(d;\psi_{\tau_c}))$. We assume $g'(d) > 0$, $\beta'_{\tau_c}(d;\psi_{\tau_c}) < 0$, and that $g(d) \to \infty$ and $\beta_{\tau_c}(d;\psi_{\tau_c}) \to 0$, as $d \to \infty$. Additionally, $\beta_{\tau_c}(0;\psi_{\tau_c}) = \beta_{\tau_c}$, $g(0) = 0$. The platform needs to decide both the number of standby cleaners to retain and the optimal discount amount for high- and low-value customers. That is,

$$
\begin{aligned}
\max_{\overline{C}\in\mathbb{Z}^+\cup\{0\}} \mathbb{E}_{\alpha,C}\Bigg[& \underbrace{k_H(1-\alpha)\gamma L + k_L(1-\alpha)(1-\gamma)L}_{\text{Value from cleanings not affected by no-show}} \\
& + \max_{0\le y\le \min\{\gamma\alpha L, C\},\, d_H\ge 0} \Bigg\{ \underbrace{k_{o,H}y}_{\text{Type-H, standby}} \\
& + \underbrace{(\gamma\alpha L - y)(v_H(1-\beta_H(d_H;\psi_H)) - g(d_H))}_{\text{Type-H, discounts}} \\
& + \max_{0\le z\le \min\{(1-\gamma)\alpha L, C-y\},\, d_L\ge 0} \Bigg[\underbrace{k_{o,L}z}_{\text{Type-L, standby}} \\
& + \underbrace{((1-\gamma)\alpha L - z)(v_L(1-\beta_L(d_L;\psi_L)) - g(d_L))}_{\text{Type-L, discounts}} \Bigg]\Bigg\}\Bigg] \\
& - \underbrace{\mu\overline{C}}_{\text{Cost of standby cleaners}},
\end{aligned} \tag{7}
$$

where $k_{\tau_c} = v_{\tau_c} + p - m$ and $k_{o,\tau_c} = k_{\tau_c} - m_o$; $y \ge 0$ and $z \ge 0$ determine the allocation of standby cleaners across customer types.

5.2.2. Solution Characterization. We solve for the optimal remedy strategy $(\overline{C}^*, y^*, z^*, d^*_{\tau_c})$ from the model above. For ease of notation, we let $s_{\tau_c}(d) := s(d;\psi_{\tau_c}) = v_{\tau_c}(1-\beta_{\tau_c}(d;\psi_{\tau_c})) - g(d)$ denote the value recovered through a frustration discount d from customer type τ_c. We assume $s_{\tau_c}(d)$ is concave in d for $d \ge 0$. Further denote $-\tau_c$ as the opposite customer type. That is, $-H = L$ and $-L = H$.

Proposition 1.

The optimal remedy strategy has the following form.

1. *The optimal solution to the maximization problem is characterized as follows.*

(a) *The optimal frustration discounts to offer are* $d^*_{\tau_c} = \arg\max_{d_{\tau_c}\ge 0} s_{\tau_c}(d_{\tau_c};\psi_{\tau_c})$.

(b) *Suppose* $k_{o,\tau_c} \ge s_{\tau_c}(d^*_{\tau_c})$ *and* $k_{o,H} - s_H(d^*_H) \ge k_{o,L} - s_L(d^*_L)$.[15] *There is an optimal* $\overline{C}^*$, *and*

$$
\begin{aligned}
\overline{C}^* = \arg\max_{\overline{C}\in\mathbb{Z}^+\cup\{0\}} \Bigg(& M_1(s_{\tau_c}(d^*_{\tau_c}))(\beta_s\overline{C} \\
& - \int_0^\infty F_\alpha((\gamma L)^{-1}x) h_{C|\overline{C}}(x)dx) \\
& + M_2(s_L(d^*_L))(\beta_s\overline{C} - \int_0^\infty F_\alpha(L^{-1}x) h_{C|\overline{C}}(x)dx) \\
& + M_3(s_{\tau_c}(d^*_{\tau_c}), \boldsymbol{\theta}) - \mu\overline{C} \Bigg),
\end{aligned}
$$

where $h_{C|\overline{C}}(x) = \mathbb{P}(C \ge x)$ *in which* $C \sim Bin(\overline{C}, \beta_s)$, *and* $M_i(\cdot)$, $i \in \{1,2,3\}$ *are functions of parameters. Furthermore, for any realization of* α *and* C, $y^* = \min\{\gamma\alpha L, C\}$ *and* $z^* = \min\{(1-\gamma)\alpha L, C - y^*\}$. *That is, high-value customers are prioritized over low-value customers in getting a standby cleaner.*[16]

2. *We conclude the following properties of the optimal solution* $(\overline{C}^*, d^*_{\tau_c})$.

(a) *Let* $\Psi_{\tau_c} = \{\psi_{\tau_c} \mid -v_{\tau_c}\beta'_{\tau_c}(0;\psi_{\tau_c}) \le g'(0)\}$. *For all* $\psi_{\tau_c} \in \Psi_{\tau_c}$, $d^*_{\tau_c} = 0$.

(b) *Given* $\psi_{-\tau_c}$, *for any sufficiently large* μ, $\exists\psi^0_{\tau_c}(\mu)$, *such that when* $\psi_{\tau_c} > \psi^0_{\tau_c}(\mu)$, $\overline{C}^* = 0$.

We relegate all proofs to Online Appendix D. The structure of the optimal remedy strategy suggested by Part 2 of the proposition follows these tradeoffs: If the marginal net benefit of offering the first unit of discount is nonpositive, no discounts should be offered. If the net benefit of keeping standby cleaners does not outweigh the net benefit of offering discounts, no standby cleaners should be retained. Although these general intuitions are clear, how these net benefits are driven by different parameter combinations is not trivial.

5.2.3. Numerical Studies: Optimal Remedy Strategies. To illustrate the actionable insights our model can provide, we conduct extensive numerical studies to generate tailored optimal remedy strategies across substantial parameter heterogeneity among service platforms. Platforms may differ in key characteristics such as customer responsiveness to discounts, customer type-mix, standby worker availability and cost, and no-show predictability. As shown in Proposition 1, all these factors, captured by parameters including ψ_{τ_c}, γ, μ, β_s, μ_α, and σ_α, collectively shape platforms' optimal decisions.

The numerical results are shown in Figure 1. Consider Figure 1(a). We fix cleaner availability at $\beta_s = 0.9$, the

Figure 1. Optimal Remedy Strategies Tailored to Customer and Platform Characteristics

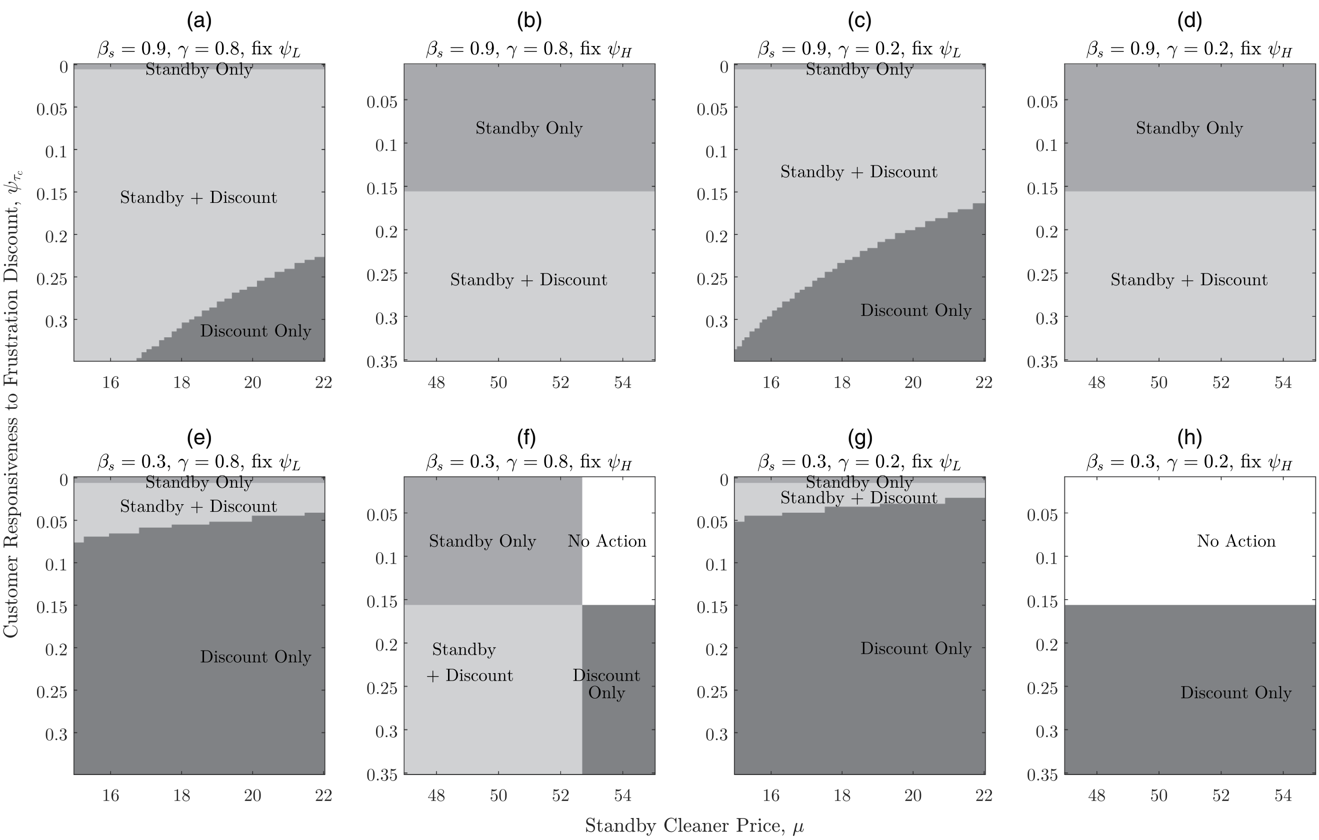

Notes. This figure plots platforms' optimal remedy strategies under varying parameters. We numerically solve model (7) to determine the optimal remedy strategy. In each subplot, we fix β_s and γ as indicated in the subplot title. Other model parameters $(L, v_H, v_L, p, m, m_o, \beta_H, \beta_L, \mu_\alpha, \sigma_\alpha)$ are calibrated from our data and empirical findings (see Online Appendix E). We assume $\alpha \sim Unif$ and $s_{\tau_c}(d) = v_{\tau_c}(1 - \beta_{\tau_c}\exp(-\psi_{\tau_c} d)) - d$. For each subplot, we fix $\psi_{-\tau_c} = 0.0005$ and vary standby cleaner price μ and customers' responsiveness to frustration discounts, ψ_{τ_c}.

share of high-value customers at $\gamma = 0.8$, and set low-value customer responsiveness to $\psi_L = 0.0005$. The panel plots a platform's optimal strategy for different values of μ and ψ_H. When ψ_H is low, frustration discounts are ineffective, and the platform relies solely on standby cleaners, as shown in Proposition 1(2(a)). As ψ_H increases, customers' responsiveness to frustration discounts increases and discounts become more cost-effective, leading to a shift toward a hybrid strategy. Eventually, beyond the threshold ψ_H^0, it becomes cheaper to forgo standby cleaners entirely and rely on post–no-show discounts, as shown in Proposition 1(2(b)). This threshold depends inversely on μ: The higher the cost of retaining standby cleaners, the less customer responsiveness is required to justify a discount-only strategy. These shifts are reflected in the transition from light to dark gray regions in the figure.

This fundamental tradeoff between discount effectiveness and standby costs manifests consistently across all parameter combinations, although with important variations in regions of optimal strategy. Comparing scenarios with high standby cleaner availability ($\beta_s = 0.9$) versus low availability ($\beta_s = 0.3$)—for instance, Figure 1(a) versus Figure 1(e), and Figure 1(b) versus Figure 1(f)—reveals that lower availability substantially diminishes the viability of the standby cleaner strategy. Specifically, the region in which the platform chooses to retain standby cleaners shrinks considerably, because uncertainty about cleaner presence renders discounts comparatively more cost-effective. Similarly, the proportion of high-value customers significantly influences optimal strategy selection. Comparing Figure 1(a) versus Figure 1(c) and Figure 1(e) versus Figure 1(g), platforms with predominantly high-value customers ($\gamma = 0.8$) should more aggressively retain standby cleaners than those with fewer of these customers ($\gamma = 0.2$). The bottom-right panel, Figure 1(h), highlights an important boundary case where neither remedy is justified. When standby labor costs are prohibitive and customers exhibit minimal discount responsiveness, the economically rational choice is to accept customer attrition without intervention.

Note that regions labeled "Standby Only" or "Standby + Discount" in Figure 1 indicate that retaining standby cleaners is optimal, but the specific optimal number to retain varies. Figure 2 illustrates the optimal number of standby cleaners under various parameter combinations. Whereas lower unit costs of standby cleaners (μ) increase

Figure 2. Optimal Number of Standby Cleaners on Standby Cleaner Availability

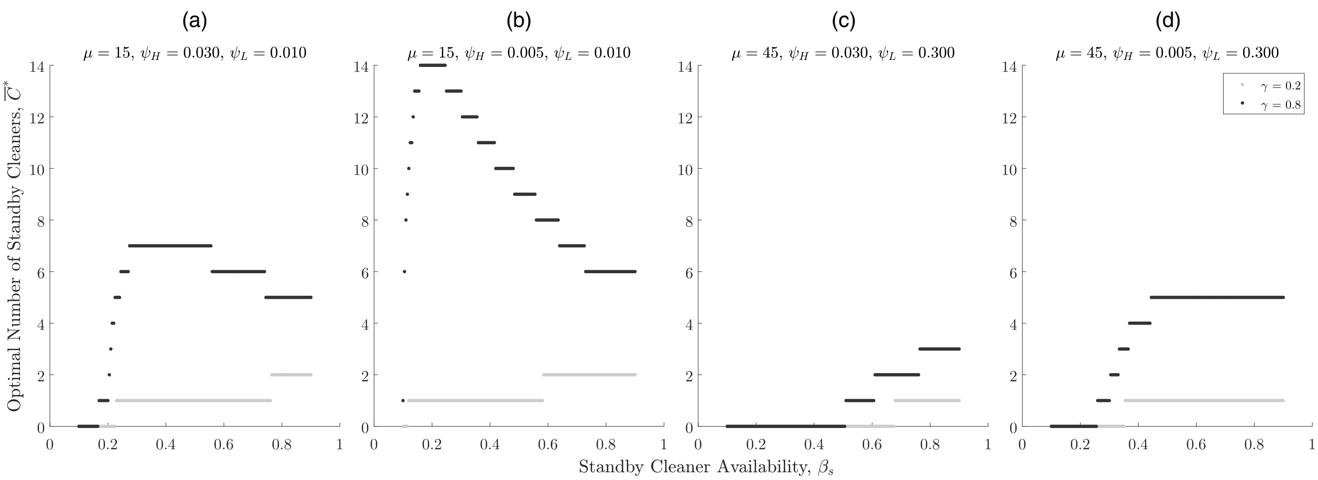

Notes. This figure shows the optimal number of standby cleaners to retain for different values of standby cleaner availability β_s. Each subplot presents results under two customer mix scenarios: $\gamma = 0.2$ (predominantly low-value customers) and $\gamma = 0.8$ (predominantly high-value customers). Aside from μ, ψ_H, and ψ_L as shown in the subplot titles, all other parameters follow the same values used in Figure 1.

optimal standby cleaner $\overline{C}^*$ as expected, the relationship between cleaner availability (β_s) and optimal standby cleaner pool size follows a notable nonmonotonic pattern. At low availability levels, investing in standby cleaners becomes inefficient due to the low probability of their actual presence, whereas at high availability levels, the marginal benefit of additional standby cleaners decreases due to reduced variability in turnout rates. These two forces result in the observed inverted U-shape. This nuanced interaction between cleaner availability and unit costs highlights the challenges platforms face when designing remedy strategies and demonstrates the practical value our analytical model can provide.

Our numerical analysis uses a single-day planning horizon and a single-city context. However, a service platform can adapt the planning horizon and geographical constraints to suit its needs. For instance, if the results from our prediction model indicate that mornings are particularly susceptible to no-shows, the platform can consider planning for standby cleaners specifically for the morning. Our model is flexible enough to accommodate various remedy planning scenarios.

Specifically, the optimal strategy for our focal platform is to keep seven standby cleaners prioritized for high-value customers, extend a 65 LC discount (54% of the per-cleaning price) to high-value customers, and issue no discounts to affected low-value customers. The parameters used for this analysis are presented in Online Appendix E. On average, compared with taking no action at all, the optimal remedy strategy allows our platform to recoup 1.2% of total customer lifetime value. The gain is not negligible, considering that only 3.51% of cleanings are affected by no-shows.

Thus far, we considered only short-term impacts, and our current model does not capture long-term platform losses from no-shows. To address this, we construct a two-stage model to examine how no-shows affect long-term platform size and revenue. Comparing scenarios with and without remedy strategies, we find that remedy strategies, such as retaining standby cleaners, can significantly mitigate compounding revenue losses. These results reinforce the importance of service failure recovery strategies for platform sustainability and growth. Detailed analyses of these long-term dynamics are provided in Online Appendix G.

6. Conclusion

No-shows are pervasive across service industries, yet their consequences remain underexplored—particularly in online-to-offline service settings where each customer is matched one-to-one with a single service provider. Drawing on granular transaction data from a European home-cleaning platform, our study fills this gap. Our findings show that service provider no-shows significantly increase customer attrition, with loyal customers exhibiting heightened sensitivity to these service failures. Beyond customer attrition impacts, no-shows substantially reduce customer lifetime value.

Our analysis further indicates that predicting no-shows is challenging, highlighting the critical importance of effective postincident remedy strategies. To this end, we examine the effectiveness of two remedy strategies and find that offering discounts can alleviate customer attrition following no-shows, but simply restoring the cleaning for the same day seems ineffective. Building on these empirical insights, we develop an analytical framework that platforms can calibrate using their own data to determine optimal remedy strategies.

The findings of our study can be naturally extended to other online-to-offline home-service platforms involving gig workers—plumbing services booked through Angi, appliance installation arranged via TaskRabbit, electrician visits scheduled on Urban Company, or home cleaning services booked through Handy. In these settings, the core ingredients remain unchanged; a gig worker is matched to one household task for a specific date and time of service, and customer decision processes are also similar to those in our study. After a customer specifies a service request, they are matched with a service provider. Importantly, the platforms do not assign service providers to service requests, and service providers have autonomy on when and where to work. On the day of the service, either the service takes place as promised, or the service provider is absent. Customers then decide whether to return to the platform for future purchases or to churn.

Because these companies operate similarly to our focal platform, they can leverage our insights to combat no-shows effectively. Our study offers a comprehensive framework that quantifies the impact of no-shows on attrition and CLV, estimates cleaning-level no-show probabilities, and finally inputs these parameters into an optimization model to design tailored remedy strategies using platforms' own data. The magnitude of these calibrated parameters might vary across different platforms. For example, service providers may exhibit different no-show rates and varying levels of no-show predictability in different settings. Customers may also respond differently to remedy strategies. Although a missed home cleaning can often be rescheduled to a later time in the day or another day with a frustration discount, a plumber's job is typically more urgent; one cannot leave a water leak unattended for hours or days. This urgency makes a plumber's no-show more disruptive and harder to offset with a simple discount or a same-day restoration, making prompt restoration through standby service providers critical. Differences in job urgency, therefore, lead to differences in the relative effectiveness of remedy strategies. On the provider side, outside options might differ: Cleaners may face a different opportunity set than handymen, resulting in variation in standby provider costs and availability. These parameter variations represent quantitative rather than qualitative differences across customers, service providers, and platforms.

Our managerial insights can be generalized to many settings, but our study is not without limitations. Ideally, a true randomized experiment—intentionally assigning some cleanings to be no-show incidents—could improve the internal validity of our estimated effects. However, deliberately degrading the service would undermine both the customer experience and the success of the platform, rendering such randomization infeasible. Nonetheless, future studies could examine a broader set of remedy strategies by randomizing restoration policies across no-show events and evaluating their effectiveness through field experiments.

In addition, implementing the optimal remedy strategies we propose might entail several operational challenges. First, ensuring a prompt dispatch of a standby service provider could be inherently difficult in gig worker settings, because these workers control their own schedules. Our model probabilistically captures the availability of standby cleaners, but real-time assignment of standby cleaners to no-show cleanings in a short window can be challenging. Second, adjusting the standby cleaner pool on a daily basis might introduce an operational burden on the platform. In this regard, platforms can implement our framework for longer time horizons. In such cases, our model can be extended to take into consideration the dynamic revelation of no-shows and the dynamic allocation of standby service providers to no-show incidents. Moreover, if platforms can predict no-shows, they could implement preemptive interventions, such as proactively contacting cleaners identified as high-risk for absence. These preventive measures can significantly reduce the costs of combating no-shows. Even so, maintaining a standby cleaner pool remains valuable because it enables prompt replacement when an absence occurs. Therefore, we propose that combining such preventative tools with mitigation strategies can help platforms develop doubly robust strategies to manage no-shows.

We hope that our study encourages future research on service failure and its recovery, especially in gig economy settings. For example, although our focal platform does not impose monetary penalties on service providers for no-shows, future research could explore the implications of such penalty mechanisms. On the one hand, penalties might serve as a preventive measure that deters absenteeism; on the other hand, they could reduce the supply of gig workers on the platform in the long run. This nuanced interaction between deterrence and supply effects presents an opportunity for future studies to examine comprehensive worker management strategies that platforms can implement in the face of no-shows. Our paper specifically focuses on service provider no-shows, but the underlying framework is generalizable to other service failures and a broader range of remedy strategies, offering a springboard for future research in this understudied area of service operations.

Acknowledgments

The authors thank the department editor, associate editor, and anonymous reviewers for their constructive feedback.

The authors are especially grateful to Marty Lariviere for his guidance and support throughout the development of this project.

Endnotes

[1] U.S. Bureau of Labor Statistics (2022) Absences from work of employed full-time wage and salary workers by occupation and industry. Accessed March 2, 2026, https://www.bls.gov/cps/aa2021/cpsaat47.htm.

[2] Throughout the paper, we use no-show and absenteeism interchangeably.

[3] Grand View Research (2022), Online On-demand Home Services Market (2022–2030). Accessed March 2, 2026, https://www.grandviewresearch.com/industry-analysis/online-on-demand-home-services-market-report.

[4] For 179 no-show cleanings, the ID of the absent cleaner is missing. Results remain robust when these cleanings are excluded.

[5] One of the observations in our sample has a missing price value and thus was dropped from our analysis.

[6] Throughout the paper, we refer to this event interchangeably as "leaving the platform" or "discontinuing use of the platform."

[7] We interpret the estimate, 0.296, in terms of the change in attrition probability. Specifically, for each observation, we predict the probability of attrition twice—once setting the no-show indicator to one and once setting it to zero—while holding all other covariates at their observed values. We then take the percentage difference between these two predicted probabilities and average this difference across all observations in the sample. This procedure directly yields the average percentage change in attrition probability associated with a no-show. Our estimate implies that experiencing a no-show is associated with a 25.6% increase in the customer attrition probability.

[8] We use Python's XGBoost package for the estimation. XGBoost has shown strong predictive performance (Clarke et al. 2020).

[9] The platform fully absorbs the cost of the frustration discount, and cleaners' incomes are not affected by the discount.

[10] For example, if cleaning i occurred on Tuesday, August 20, 2019, in city A1 of country A, the instrument captures the average frustration discount issued in country A on all Tuesdays in August 2019, excluding city A1 on August 20, 2019.

[11] The *SameDayRestored* remedy is only relevant in case of no-shows. Therefore, it is perfectly collinear with the *SameDayRestored* × *NoShow*, and we cannot separately identify this interaction term.

[12] Each cleaning is associated with a predicted no-show probability, $\hat{y}_i$, from the XGBoost prediction model. Assuming conditional independence of no-shows across cleanings (after controlling for month-year × city fixed effects), the expected no-show rate is $L^{-1}\sum_{i=1}^{L}\hat{y}_i$, and the standard deviation is $\sqrt{L^{-2}\sum_{i=1}^{L}\hat{y}_i(1-\hat{y}_i)}$.

[13] The high lifetime value customers should reliably align with more loyal customers of the platform.

[14] In this specification, because v_{τ_c} is the customer lifetime value, we implicitly assume that the baseline attrition rates of high and low-value customers are already accounted for in the lifetime value v_{τ_c}. Therefore, the interpretation of β_{τ_c} is the additional attrition rate among type τ_c customers induced by no-shows.

[15] We present the proposition and numerical analysis in the main text based on conditions $k_{o,\tau_c} \geq s_{\tau_c}(d^*_{\tau_c})$ and $k_{o,H} - s_H(d^*_H) \geq k_{o,L} - s_L(d^*_L)$ because these conditions are satisfied by parameters calibrated from our data.

[16] Note that our model (7) does not assume the platform always gives high-value customers priority access to standby cleaners. When $k_{o,\tau_c} \geq s_{\tau_c}(d^*_{\tau_c})$ and $k_{o,H} - s_H(d^*_H) < k_{o,L} - s_L(d^*_L)$, the platform reverses prioritization; low-value customers receive standby cleaners first.

References

Allon G, Cohen MC, Sinchaisri WP (2023) The impact of behavioral and economic drivers on gig economy workers. *Manufacturing Service Oper. Management* 25(4):1376–1393.

Allon G, Federgruen A, Pierson M (2011) How much is a reduction of your customers' wait worth? An empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing Service Oper. Management* 13(4):489–507.

Anderson RE (1973) Consumer dissatisfaction: The effect of disconfirmed expectancy on perceived product performance. *J. Marketing Res.* 10(1):38–44.

Anderson SW, Baggett LS, Widener SK (2008) The impact of service operations failures on customer satisfaction: Evidence on how failures and their source affect what matters to customers. *Manufacturing Service Oper. Management* 11(1):52–69.

Anderson ET, Fitzsimons GJ, Simester D (2006) Measuring and mitigating the costs of stockouts. *Management Sci.* 52(11): 1751–1763.

Bachmann P, Meierer M, Näf J (2021) The role of time-varying contextual factors in latent attrition models for customer base analysis. *Marketing Sci.* 40(4):783–809.

Clarke J, Chen H, Du D, Hu YJ (2020) Fake news, investor attention, and market reaction. *Inform. Systems Res.* 32(1):35–52.

Cohen MC, Fiszer MD, Kim BJ (2021) Frustration-based promotions: Field experiments in ride-sharing. *Management Sci.* 68(4): 2432–2464.

Colombo R, Jiang W (1999) A stochastic RFM model. *J. Interactive Marketing* 13(3):2–12.

Craighead CW, Karwan KR, Miller JL (2004) The effects of severity of failure and customer loyalty on service recovery strategies. *Production Oper. Management* 13(4):307–321.

Cui R, Lu Z, Sun T, Golden JM (2023) Sooner or later? Promising delivery speed in online retail. *Manufacturing Service Oper. Management* 26(1):233–251.

Fader PS, Hardie BG, Lee KL (2005) RFM and CLV: Using iso-value curves for customer base analysis. *J. Marketing Res.* 42(4):415–430.

Fry MJ, Magazine MJ, Rao US (2006) Firefighter staffing including temporary absences and wastage. *Oper. Res.* 54(2):353–365.

Gallo A (2014) The value of keeping the right customers. *Harvard Bus. Rev.* (October 29), https://hbr.org/2014/10/the-value-of-keeping-the-right-customers.

Green LV, Savin S, Savva N (2013) "Nursevendor problem": Personnel staffing in the presence of endogenous absenteeism. *Management Sci.* 59(10):2237–2256.

Grégoire Y, Fisher RJ (2008) Customer betrayal and retaliation: When your best customers become your worst enemies. *J. Acad. Marketing Sci.* 36:247–261.

Hall J, Porteus E (2000) Customer service competition in capacitated systems. *Manufacturing Service Oper. Management* 2(2):144–165.

Hassin R, Mendel S (2008) Scheduling arrivals to queues: A single-server model with no-shows. *Management Sci.* 54(3):565–572.

Hausman JA (1996) Valuation of new goods under perfect and imperfect competition. Bresnahan TF, Gordon RJ, eds. *The Economics of New Goods* (University of Chicago Press, Chicago), 207–248.

Hess RL, Ganesan S, Klein NM (2003) Service failure and recovery: The impact of relationship factors on customer satisfaction. *J. Acad. Marketing Sci.* 31(2):127–145.

Kamalahmadi M, Yu Q, Zhou YP (2021) Call to duty: Just-in-time scheduling in a restaurant chain. *Management Sci.* 67(11): 6751–6781.

Kesavan S, Lambert SJ, Williams JC, Pendem PK (2022) Doing well by doing good: Improving retail store performance with responsible scheduling practices at the Gap, Inc. *Management Sci.* 68(11):7818–7836.

Kong Q, Li S, Liu N, Teo CP, Yan Z (2020) Appointment scheduling under time-dependent patient no-show behavior. *Management Sci.* 66(8):3480–3500.

Kwon C, Raman A (2023) The effects of inconsistent work schedules on employee lateness and absenteeism. Preprint, submitted November 2, http://dx.doi.org/10.2139/ssrn.4592763.

Li X, Ma B, Zhou C (2017) Effects of customer loyalty on customer entitlement and voiced complaints. *Service Industries J.* 37(13–14):858–874.

Liu J, Xie J, Yang KK, Zheng Z (2019) Effects of rescheduling on patient no-show behavior in outpatient clinics. *Manufacturing Service Oper. Management* 21(4):780–797.

Lu G, Du RY, Peng X (2022) The impact of schedule consistency on shift worker productivity: An empirical investigation. *Manufacturing Service Oper. Management* 24(5):2780–2796.

Mani V, Kesavan S, Swaminathan JM (2015) Estimating the impact of understaffing on sales and profitability in retail stores. *Production Oper. Management* 24(2):201–218.

McVicar A (2003) Workplace stress in nursing: A literature review. *J. Advanced Nursing* 44(6):633–642.

Nisbett RE, Wilson TD (1977) The halo effect: Evidence for unconscious alteration of judgments. *J. Personality Soc. Psych.* 35(4):250–256.

Noone BM, Lee CH (2011) Hotel overbooking: The effect of overcompensation on customers' reactions to denied service. *J. Hospitality Tourism Res.* 35(3):334–357.

Oblander ES, McCarthy D (2021) How has COVID-19 impacted customer relationship dynamics at restaurant food delivery businesses. *Marketing Sci. Institute Working Paper Ser.* 23(2139):35.

Oliva R, Sterman JD (2001) Cutting corners and working overtime: Quality erosion in the service industry. *Management Sci.* 47(7):894–914.

Oliver RL (1980) A cognitive model of the antecedents and consequences of satisfaction decisions. *J. Marketing Res.* 17(4):460–469.

Oliver RL, DeSarbo WS (1988) Response determinants in satisfaction judgments. *J. Consumer Res.* 14(4):495–507.

Parasuraman A, Zeithaml VA, Berry LL (1988) SERVQUAL: A multiple-item scale for measuring consumer perceptions of service quality. *J. Retailing* 64(1):12–40.

Pendem PK, Green PI Jr, Staats BR (2022) The microstructure of work: Understanding productivity benefits and costs of interruptions. *Manufacturing Service Oper. Management* 24(4):2202–2220.

Roy D, Mishra A, Sinha KK (2021) Taxing the taxpayers: An empirical investigation of the drivers of baseline changes in US federal government technology programs. *Manufacturing Service Oper. Management* 24(1):370–391.

Schmittlein DC, Morrison DG, Colombo R (1987) Counting your customers: Who-are they and what will they do next? *Management Sci.* 33(1):1–24.

Sivakumar K, Li M, Dong B (2014) Service quality: The impact of frequency, timing, proximity, and sequence of failures and delights. *J. Marketing* 78(1):41–58.

Tax SS, Brown SW, Chandrashekaran M (1998) Customer evaluations of service complaint experiences: Implications for relationship marketing. *J. Marketing* 62(2):60–76.

Terza JV, Basu A, Rathouz PJ (2008) Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling. *J. Health Econom.* 27(3):531–543.

Unruh L, Joseph L, Strickland M (2007) Nurse absenteeism and workload: Negative effect on restraint use, incident reports and mortality. *J. Advanced Nursing* 60(6):673–681.

Wang WY, Gupta D (2014) Nurse absenteeism and staffing strategies for hospital inpatient units. *Manufacturing Service Oper. Management* 16(3):439–454.

Wangenheim F, Bayón T (2007) Behavioral consequences of overbooking service capacity. *J. Marketing* 71(4):36–47.

Wetzel HA, Hammerschmidt M, Zablah AR (2014) Gratitude versus entitlement: A dual process model of the profitability implications of customer prioritization. *J. Marketing* 78(2): 1–19.

Wooldridge JM (2010) *Econometric Analysis of Cross Section and Panel Data* (MIT Press, Cambridge, MA).

Xu X, Liu W, Gursoy D (2019) The impacts of service failure and recovery efforts on airline customers' emotions and satisfaction. *J. Travel Res.* 58(6):1034–1051.

Xu Y, Lu B, Ghose A, Dai H, Zhou W (2023) The interplay of earnings, ratings, and penalties on sharing platforms: An empirical investigation. *Management Sci.* 69(10):6128–6146.

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

Zeithaml VA, Berry LL, Parasuraman A (1996) The behavioral consequences of service quality. *J. Marketing* 60(2):31–46.