The Business School for the World®

Working Paper

2023/64/TOM

(Revised version of 2020/35/TOM)

Select, Swipe, and Serve: Examining the Impact of Food-Delivery Platforms on Restaurant Demand Characteristics

Varun Karamshetty
INSEAD, varun.karamshetty@insead.edu

Michael Freeman
INSEAD, michael.freeman@insead.edu

Sameer Hasija
INSEAD, sameer.hasija@insead.edu

Last updated on November 17, 2023

Problem Definition: Restaurants increasingly depend on food-delivery platforms to cater to consumer demand. To devise an effective platform strategy, restaurants must understand how these platforms change the characteristics of their demand and the factors that moderate this change.
Methodology and Results: Using a unique transaction-level dataset of over 50 million transactions from a quick service restaurant chain comprised of 99 restaurants, we assess the effect of platform dependence on restaurant demand characteristics, focusing on order composition and demand forecast accuracy. Our analysis reveals that a ten-percentage point increase in delivery platform dependence leads to a 10.5% reduction in the proportion of sales revenue from high-margin items and a 2.83% rise in demand forecast error. Moreover, we show that the influence on high-margin sales is moderated by the following customer characteristics: affluence, group size, and order timing.
Managerial Implications: These findings underscore hidden costs associated with platform dependence, extending beyond commission fees to include operational challenges and profit margin reductions. For practitioners, our research prompts restaurants to re-examine their digital strategies, balancing the benefits of platform reach with the need for optimized operational efficiency and profitability. For academics, this study paves the way for further explorations into the dynamic relationship between restaurants and digital platforms, a critical aspect of the modern food service industry.

Keywords: Digital Platforms; Demand Characteristics; Operational Efficiency; Panel Analysis

Select, Swipe, and Serve: Examining the Impact of Food-Delivery Platforms on Restaurant Demand Characteristics

Varun Karamshetty
National University of Singapore, Singapore, varun.karamshetty@nus.edu.sg

Michael Freeman
INSEAD, Singapore, michael.freeman@insead.edu

Sameer Hasija
INSEAD, Singapore, sameer.hasija@insead.edu.

Problem Definition: Restaurants increasingly depend on food-delivery platforms to cater to consumer demand. To devise an effective platform strategy, restaurants must understand how these platforms change the characteristics of their demand and the factors that moderate this change. **Methodology and Results**: Using a unique transaction-level dataset of over 50 million transactions from a quick service restaurant chain comprised of 99 restaurants, we assess the effect of platform dependence on restaurant demand characteristics, focusing on order composition and demand forecast accuracy. Our analysis reveals that a ten-percentage-point increase in delivery platform dependence leads to a 10.5% reduction in the proportion of sales revenue from high-margin items and a 2.83% rise in demand forecast error. Moreover, we show that the influence on high-margin sales is moderated by the following customer characteristics: affluence, group size, and order timing. **Managerial Implications**: These findings underscore hidden costs associated with platform dependence, extending beyond commission fees to include operational challenges and profit margin reductions. For practitioners, our research prompts restaurants to re-examine their digital strategies, balancing the benefits of platform reach with the need for optimized operational efficiency and profitability. For academics, this study paves the way for further explorations into the dynamic relationship between restaurants and digital platforms, a critical aspect of the modern food service industry.

Key words: digital platforms; demand characteristics; operational efficiency; panel analysis

History: Last updated on November 17, 2023

---

1. Introduction

The last decade has seen a substantial increase in the prevalence and usage of food-delivery platforms, largely driven by consumers' growing demand for convenience. This trend has seen the global food-delivery market expand significantly, reaching an estimated value of US$150 billion in 2021, almost three times its value in 2017 (Ahuja et al. 2021). Reliance on these platforms has become even more pronounced since the COVID-19 pandemic, as they provided a means for many restaurants to maintain operations amidst the ensuing difficulties.

Food-delivery platforms are widely recognized as an efficient avenue for restaurants to broaden their reach and enhance their customer base (Sharma and Mehrotra 2007, Xia and Zhang 2010). However, their impact on restaurant profitability remains a complex and contested issue (Houck 2017, Dunn 2018). Contemporary debate revolves around the high commission fees imposed by these platforms and the potential for cannibalization of traditional in-store sales (Forman 2019, Haddon and Jargon 2019). This paper expands the discourse beyond these two issues, specifically, by exploring how digital food-delivery platforms affect the sales of high-margin items and the accuracy of demand forecasting. Both these demand characteristics are crucial elements in the restaurant industry, where the competitive edge depends on the sale of high-margin supplementary items, such as beverages, desserts, and sides (Ransom 2009, Schlosser 2012), combined with operational efficiencies realized through standardization and accurate demand prediction (Wiggers 2018, Magnin 2019).

Extant theory and anecdotal evidence lead to ambiguous conclusions about how food-delivery platforms impact these demand characteristics. Behavioral studies suggest online customers may be less price-sensitive and more willing to purchase high-margin items (Andrews and Currim 2004) whereas anecdotal evidence suggests delivery orders may contain fewer of these items due to quality concerns (Haddon and Jargon 2019) and reduced switching costs, which might motivate customers to buy high-margin items from specialized establishments instead. Meanwhile, food-delivery platforms could either help or hinder restaurants in terms of demand forecasting. These platforms can expand a restaurant's reach, potentially reducing forecasting error via demand pooling, but as customers can easily switch between restaurants, competition increases, and demand on any given day is possibly more unpredictable (Mahajan 2019). Overall, we lack sufficient understanding of the influence of food-delivery platforms on these vital demand characteristics.

This investigation uses a novel, proprietary database of detailed transaction-level data from a quick service restaurant (QSR) chain covering every transaction (approximately 50 million) made across 99 outlets in 2018. For each transaction, we observe the stock-keeping unit (SKU), number of units sold, timestamp, and delivery method (i.e., dine-in, take-away, or made through a food-delivery platform). Using this dataset, we calculate the sales made through the platform as a proportion of total sales for each restaurant-day, forming the basis for our primary independent variable, which we term "platform dependence". In our dataset, the median platform dependence grows from 6.1% at the start of 2018 to 12.9% by year-end, with significant variability throughout (see Fig. 1a in Appendix A).

We apply panel analysis techniques to take advantage of this within-restaurant variation in dependence and identify that platform dependence significantly impacts both the demand characteristics considered. Specifically, we find that a ten percentage-point (pp) increase in a restaurant's dependence on delivery platforms is associated with a 10.5% reduction in the proportion of sales revenue from high-margin items. Additionally, we observe a 2.83% increase in the demand forecast error, which was measured as the disparity between forecasted and observed sales.

These findings have considerable implications for the profitability of restaurants. Beyond the commissions paid to the food-delivery platforms, our results suggests that two principal cost components in the restaurant industry – ingredients and staffing – are susceptible to inflation under a platform-dependent business model. Ingredient costs are inextricably tied to both the composition of orders and the level of demand uncertainty: operations management (OM) theory robustly links higher demand uncertainty with an increase in food wastage, resulting in higher ingredient costs; concurrently, an increase in the proportion of sales from low-margin items corresponds to a relatively higher ingredient and production costs. Staffing costs may also rise, as when demand uncertainty is high, a larger buffer staff is needed, leading to lower staff utilization. It is therefore important for restaurants to factor in these potential changes to their overall cost structure when contracting with food-delivery platforms.

Nevertheless, as food-delivery platforms have become an inescapable fixture in the market, it is important to understand the circumstances under which these adverse effects are diminished and intensified. A review of the rich literature on restaurant management indicates three factors that could moderate the relationship between platform dependence and the demand characteristics considered in our paper (Thompson 2010). Our subsequent analysis of the moderating effects of these factors on the relationship between platform dependence and high-margin item sales reveals that the effects are more pronounced in restaurants with a higher peak-demand at lunchtime (such as those catering to office workers) than those with a higher peak-demand at dinnertime. On the other hand, the impact is milder on restaurants frequented by larger groups (such as families) than those visited by smaller groups. Similarly, the impact is milder on restaurants with a higher spend on core menu items (e.g., burgers). None of these factors, however, was found to moderate the relationship between platform dependence and forecast accuracy.

These findings have implications for both theoretical and empirical OM studies. An established body of OM research studies the downstream effects of demand characteristics – for instance, the impact of demand uncertainty on waiting times (Allon et al. 2011) – but our work redirects attention towards the factors shaping these characteristics themselves. We reveal that these fundamental demand characteristics are malleable and can shift due to platform dependence, providing an essential building block for future exploration. Furthermore, our work extends the discourse on digital platforms beyond the conventional focus on commission fees and sales cannibalization. Future research can build on our findings when modeling the complex interplay between restaurants and food-delivery platforms to give a more comprehensive understanding of platform dependence, its determinants, and its extensive implications.

2. Literature Review

There is growing interest in understanding the strategies that underpin business models in the app-based sharing economy and for digital platforms. Our paper contributes to three inter-linked literature streams. In the sections below, we give an overview of these streams and position our research within this context.

2.1. Digital Platforms and their Impact

OM researchers have long studied various aspects of digital platforms, including their foundational theories, optimization of the underlying systems and technologies, paths to adoption by participants, and the operational implications of participation.

The foundational theory of digital platforms and electronic commerce, more generally, attracted significant attention even before the recent advances in digital technology made such platforms ubiquitous. Early work by Malone et al. (1987) predicted that advances in technology would shift more economic activity onto digital platforms, and multiple studies have since modeled how such platforms can reduce transaction costs (Bakos 1997). The resulting reduced intermediation costs and more efficient inter-organizational transactions (Bakos 1991) enabled a movement towards "frictionless" markets (Bakos 1998). This theoretical foundation was later supported by empirical studies that not only tested the theoretical hypotheses but also added nuance to our understanding of the theory of digital platforms (e.g. Smith et al. 1999, Brynjolfsson and Smith 2000). A comprehensive review of the literature on the theory of digital platforms is outside the scope of this paper; we refer the reader to Wang et al. (2008) and Standing et al. (2010) for a deeper dive.

More recently, as economic activity in different sectors has indeed shifted to digital platforms (such as ride-sharing, food-delivery, and gig-working), researchers investigated how to optimize the underlying systems of these platforms by looking at interconnected aspects of platform system design, such as auction types and strategies (e.g., Bapna et al. 2004, Mithas and Jones 2007, Bapna et al. 2009), procurement strategies (e.g., Chen et al. 2005, Chandrashekar et al. 2007), pricing mechanisms (e.g., Banerjee et al. 2016, Cachon et al. 2017) and order dispatching protocols (e.g., Chen et al. 2019, Lyu et al. 2019, Özkan and Ward 2020). Despite this wealth of research on system design, however, few studies have focused on the implications for platform participants (Kapoor and Agarwal 2017). Indeed, in their review of the literature on digital platforms, Standing et al. (2010) identified that too much emphasis on digital platforms and their workings had left a research gap in terms of the impact of such business models on participating firms. While some scholars have evaluated these business models from a participant-centric strategic lens (e.g., Soh et al. 2006) and others have examined the associated benefits and costs (e.g., Standing et al. 2006), the operational implications of digital platforms for participating firms remain largely understudied.

There are emerging perspectives in other academic communities, such as Marketing and Information Systems, that investigate the repercussions of digital platform engagements on established businesses. For example, Zervas et al. (2017) and Chang and Sokol (2022) study the impact of Airbnb on the hotel industry, Cramer and Krueger (2016) and Nie (2017) analyze the impact of ride-hailing platforms on the transportation industry, and Ryu et al. (2022) explore the impact of crowdfunding platforms on the venture capital business. Our paper contributes to this developing conversation from an OM perspective by spotlighting the implications of digital food-delivery platforms on the food and beverages industry.

Moreover, our work complements recent OM studies that use analytical and empirical models to study how food-delivery platforms affect demand characteristics at a restaurant. For instance, recent theoretical models suggest that food-delivery platforms may expand a restaurant's customer base but also cannibalize its traditional demand channels (Chen et al. 2022, Feldman et al. 2023). Li and Wang (2020) empirically investigate the substitution effect of food-delivery platforms on restaurants' takeout and dine-in channels and the net impact on restaurant revenue. Meanwhile, Cui et al. (2022) investigate the impact of food-delivery platforms on mitigating revenue loss for small-business restaurants during the COVID-19 pandemic. In our paper, we instead study the impact of platform engagement on forecast accuracy and high-margin item sales for restaurants, as both these factors that can dramatically impact a restaurant's bottom line.

2.2. Multichannel Operations

Our paper also complements the literature on multichannel retail businesses. When a restaurant chooses to sign up on a digital platform, it opens a new sales channel and thus becomes a multichannel sales enterprise. Traditionally, OM research of multichannel systems aimed to maximize the efficiency of inventory procurement, warehousing, and distribution to satisfy demand in multiple channels (we refer readers to Burt and Sparks (2003) for a review of the relevant literature). Recently, the ubiquity of low-cost, high-speed mobile data and high-performance processors has allowed customers to switch between the online and offline channels seamlessly, thus blurring the lines between the two. This development has led to research on two fronts: firstly, how do processes in one channel impact processes in another and secondly, how customer behavior differs depending on the channel.

Concerning the first question, which deals with cross-channel spillovers, Bell et al. (2018) and Samuel et al. (2020) study how physical showrooms impact the operational efficiency and competitive landscape of online-first retailers. Meanwhile, Kumar et al. (2019) study how facilitating returns in an offline channel affects sales on the online channel. Although this literature is relevant to our study, the restaurant industry differs from traditional retail in an important way: the finished goods (cooked food items, in our context) are highly perishable. Consequently, the demand from the online channel (food-delivery platform) cannot be pooled across stores and satisfied via drop-shipping or a "buy-online, ship-to-store" policy. Moreover, temporal and geographical demand pooling might suit online grocery retailers (cf. Belavina et al. 2017, Astashkina et al. 2019), but the short shelf-life of finished goods (and the impatience of hungry customers) compels restaurants to centralize all production at the store level. Since restaurants regularly deal with uncertain demand for perishable dishes, they rely heavily on demand forecasts to optimize their operations and must minimize waste while maintaining service quality. Our study contributes to this important conversation by investigating how third-party food-delivery platforms are impacting restaurants' ability to accurately forecast their demand.

On the second question, which addresses differences in behavior across channels, previous research has shown that the behaviour of online and offline customers behave differently, but the differences are context-dependent and not easily generalized. For instance, consumers shopping for books and CDs consumers appear to be more price-sensitive online than offline (Brynjolfsson and Smith 2000), whereas the reverse is true for consumers shopping for groceries (Andrews and Currim 2004). Our study complements this line of research by investigating how the sale of high-margin ancillary items differs in the platform vs. traditional sales channels.

2.3. Restaurant Operations Management

As innovative restaurant technologies and analytical tools have emerged, OM researchers have used them to gain deeper insights into restaurant operations management by investigating factors that affect customer behavior and demand characteristics, both of which are critical for optimizing a wide variety of operational parameters in a restaurant Roy et al. (2022), Thompson (2010). For instance, using data from the QSR industry, Allon et al. (2011) find that customers attribute a very high cost to the time they spend waiting and will trade off waiting time with price. Similarly, Tan and Netessine (2020) study of the impact of a novel customer-facing table-top technology on restaurant demand characteristics indicates a likely increase in average sales per check as well as a shorter meal duration.

We add to this conversation by using our extensive, proprietary database to study the impact of food-delivery platforms on the sale of ancillary high-margin items (also known as *cross-selling*) and demand forecasting accuracy. The restaurant industry has razor-thin margins, and the composition of sales has a significant impact on restaurant profits. Forecast accuracy also plays an important role in reducing a restaurant's food wastage (Pereira 2018) and can help regulate customers' waiting time, which, in turn affects consumption (Ülkü et al. 2020) and restaurant revenues (De Vries et al. 2018). However, to the best of our knowledge, ours is the first study to investigate, empirically or otherwise, how food-delivery platforms are impacting traditional restaurants' ability to promote these high-margin items and accurately forecast their demand.

3. Restaurant Context

The restaurant industry operates on notoriously thin profit margins, typically averaging around 4-5% of revenue (Lunden 2020). Within this challenging landscape, QSRs, the source of our data, have carved out a distinct niche. Focusing on affordability and speed of service, QSRs pair low prices on core menu items such as burgers and chicken wings with operational efficiencies to build profitability and a unique value proposition in the market.

This business model leverages two key strategies to boost revenues while minimizing costs. On the revenue side, QSRs rely heavily on upselling and cross-selling high-margin supplemental items such as drinks, desserts, and sides, which have substantially higher margins than core menu items. These items, often sold through bundle deals, both enhance the customer experience and can significantly improve profit margins (Ransom 2009, Schlosser 2012). Various tactics are employed to encourage purchase, including suggestive selling by staff and combo meal deals. Simultaneously, QSRs strictly regulate costs through precise demand forecasting. While forecasting is important across the restaurant industry, it is crucial for QSRs due to their business model. As committing to rapid service requires streamlined processes that limit flexibility to adapt in real-time, accurate demand projections enable the optimization of ingredient quantities, staff schedules, and item preparation (Wiggers 2018, Magnin 2019). Without precise forecasts, QSRs face considerable inefficiencies through overstaffing, food wastage, and lost sales that rapidly deflate profit margins (Cachon and Terwiesch 2016, p. 390-391).

In particular, to enable the necessary speed of service, QSR operations differ considerably from those of casual dining establishments. QSRs rely on long-shelf-life frozen ingredients, such as pre-formed burger patties and French fries, that undergo initial cooking ahead of real-time demand, with the exact quantities determined by demand forecasts. The cooked food is then held in specialized hot-holding ovens engineered to maintain quality and safety during storage (Shufflebotham 2022). Strict protocols govern the hot-holding process, tracking duration and mandating the disposal of food exceeding its allocated holding time. When customers place orders, pre-cooked ingredients are pulled from the hot-holding ovens and rapidly assembled to to fulfill the order. This process preserves consistency in taste, texture, and appearance, while satisfying expectations of speedy service. (Readers can see a detailed description of this process on the McDonald's official website: `https://mcdonalds.com.lb/en/questions/question/what-happen-to-the-food-which-was-not-purchased-at-the-end-of-the-day`.

Decoupling cooking from real-time demand necessitates accurate forecasting to align projected and actual sales. QSRs generate hourly demand forecasts from historical data patterns, future events, and managerial judgment. However, errors are inevitable and miscalculations lead to substantial inefficiencies including overstaffing, food wastage, and lost sales that undermine profitability (Cachon and Terwiesch 2016, p. 390-391). Precise forecasts ensure that ingredient quantities match anticipated demand across all ordering channels. This operational approach differs considerably from that of casual dining establishments, where some post-order preparation is standard. Versatile pre-made bases that are completed to order – a technique referred to as postponement in the OM literature (Swaminathan and Lee 2003) – involving additional cooking steps that contribute to an expected longer waiting time than experience in QSRs. Ultimately, therefore, the coordinated processes centered around forecasting enable QSRs to deliver affordable items rapidly while assuring consistent quality, but the model is reliant on forecast accuracy to minimize inherent risks and sustain profit margins.

4. Hypothesis Development

We now turn to developing hypotheses regarding the impact of third-party food-delivery platforms on forecast errors and the sale of high-margin items.

4.1. Forecast Errors

By joining a third-party food-delivery platform, a traditional restaurant adds an online sales channel to its existing offline one, allowing customers to switch between the two as they prefer. Food-delivery platforms can extend a restaurant's reach (Sharma and Mehrotra 2007) and increase its market size (Xia and Zhang 2010). A larger pool of potential customers could lead to lower relative errors in forecasting demand, a phenomenon well-documented in the risk pooling literature (cf. Simchi-Levi et al. 2003). Furthermore, factors such as sudden weather changes or traffic conditions, which could deter customers from physically visiting a restaurant, do not affect the online sales channel. In fact, demand from the online and offline channels may negatively correlate, reducing overall demand variability. Thus, as restaurants increasingly rely on delivery platforms, demand patterns should become easier to forecast. Thus, we hypothesize that:

Hypothesis 1a: *Increased dependence on third-party food-delivery platforms increases the accuracy of demand forecasts*

On the other hand, food-delivery platforms can create a more challenging competitive environment for a restaurant. Any restaurant can join these platforms, including virtual kitchens that have no physical presence (cf. Isaac and Yaffe-Bellany 2019). Therefore, restaurants on these platforms are competing not just with local establishments but with any restaurant that delivers to the same area. This situation could mean more competition, including from places that were not competitors offline. Additionally, the ease of comparing prices, ratings, and menus on these platforms could lead to increased price sensitivity among customers (Lynch Jr and Ariely 2000), making daily demand more volatile and harder to predict.

Online channels can also change customers' ordering patterns. Offline customers may visit restaurants at predictable times due to socializing or workplace restrictions. However, delivery platforms remove these constraints, allowing customers to order whenever they prefer. Again, this possibility could make daily demand harder to predict and increase forecast errors. Thus, we hypothesize:

Hypothesis 1b: *Increased dependence on third-party food-delivery platforms decreases the accuracy of demand forecasts*

Given these competing arguments, it is unclear how third-party food-delivery platforms will affect a restaurant's ability to accurately forecast demand. In this paper, we seek to resolve this dilemma empirically by using our novel, proprietary transaction-level dataset.

4.2. High-margin sales

As noted in earlier sections, high-margin ancillary items like drinks, desserts, and side dishes hold the highest margins and are key contributors to a restaurant's bottom line (Ransom 2009, Schlosser 2012). Restaurants often push these high-margin items to customers through strategies such as bundled deals and up-selling. Online platforms may offer restaurants a more effective way to promote the sales of these profitable offerings through product pictures, detailed descriptions, and featured customer reviews. Furthermore, the easy-to-use interface on food-delivery platforms can encourage impulse buying of high-margin items (Parboteeah et al. 2009).

Behavioral research comparing online (e.g., platform sales) and offline customers has shown that online customers are less price sensitive, which might imply a greater willingness to spend on high-margin items (Andrews and Currim 2004). OM research has also shown that deploying self-order technologies even in an offline setting is associated with ordering more high-margin items, such as alcoholic drinks (Tan and Netessine 2020). These insights suggest that the food-delivery platforms might help to enhance the sales of high-margin items at restaurants. This leads us to hypothesize:

Hypothesis 2a: *Increased dependence on third-party food-delivery platforms increases the sale of high-margin items*

On the other hand, anecdotal evidence in popular media suggests that delivery orders might include fewer high-margin items (Haddon and Jargon 2019), possibly due to features specific to the restaurant industry. For instance, as products sold at restaurants have an extremely short shelf-life, customers have complained of a decline in the quality of delivered items such as fries and desserts (Siddharth Cavale 2021, Dewey 2017). Additionally, as online platforms reduce switching costs, customers may find it easier to purchase high-margin items from specialized establishments or simply consume what is on hand at home or work.

Moreover, the conventional method of in-restaurant cross-selling, where servers suggest add-ons or specials, may not effectively translate to the impersonal digital space of food-delivery platforms. In the absence of this spontaneous, interpersonal interaction, there could be reduced sales of high-margin items. Together these factors suggest that food-delivery platforms may inhibit high-margin-item sales. This leads us to hypothesize:

Hypothesis 2b: *Increased dependence on third-party food-delivery platforms decreases the sale of high-margin items*

In light of these competing arguments, the net impact of increasing dependence on third-party food-delivery platforms on the sales of high-margin items is unclear. We therefore adopt an empirical approach to identify the directionality of this effect.

5. Data Description and Variable Construction

5.1. Empirical Setting

We test the hypotheses laid out in Section 4 by taking advantage of a proprietary point-of-sale database obtained from a chain of QSRs. The database consists of detailed transaction-level data for approximately 50 million orders placed in 99 QSRs spread across the country in 2018. (The location of the QSRs is not disclosed to preserve anonymity.) Each QSR was operational for an average of 331 days in 2018, yielding 32,695 observations at the restaurant-day level. For each transaction, the database contains information on the quantity and selling price for each SKU within that transaction. Any bundles or promotions and discounts associated with each transaction are also recorded, along with a timestamp identifying when the payment was made. The database also contains information on the mode of delivery to customer for each order: dine-in, take-away, drive-through, or via a delivery platform.

In 2018, the food-delivery business in the country was dominated by a single homegrown platform which, according to industry experts, commanded a market share of approximately 85%. This market leader also accounts for 96.1% of all platform-based orders in our dataset, and 10.7% of all orders filled by the QSR chain. Despite the dominance of this one platform, another smaller platform also operated at the time our data was collected, which we also account for when measuring platform dependence.

Overall, then, 11.1% of total orders across all restaurants in our dataset were placed through a food-delivery platform. We use these orders to generate our measure of each QSR's dependence on third-party food-delivery platforms (cf. Section 5.5). We refer to the other 88.9% of the orders as *non-platform demand*, which is distributed as follows: 51.1% dine-in, 27.8% take-away, and 6.1% drive-through (numbers reported as a percentage of total demand, inclusive of platform orders). Furthermore, a small portion of demand (3.9%) is satisfied by the chain's own food-delivery service.

5.2. Forecasting Demand

In our setting 486,211 hourly forecasts were made across all restaurants in the QSR chain over the one-year sample period. Hourly forecasts are only made for menu items with short shelf-life – this includes all the core menu items (such as burgers and chicken wings) and side-dishes (such as fries), but not desserts and beverages. Managers of individual QSRs generate these hourly forecasts at the start of each working day and base them on a combination of historical sales data and the manager's own experience. These hourly forecasts are made for aggregate demand across all delivery channels since each restaurant has only one kitchen to satisfy all demand.

In our setting, the QSR managers use *sales revenue from core menu items before discounts* (denoted by $Sales_{id}$) as the metric for measuring demand. Doing so makes sense because in this context, the sales price for each core SKU is highly correlated with the ingredient cost and the labor costs involved in preparation. Moreover, QSRs do not directly track the number of diners they serve. QSR managers therefore base their inventory and staffing decisions on the forecasted sales value in that hour rather than the numbers of diners they expect to see. We therefore use $Sales_{id}$ as the primary metric to measure demand in this paper. Figure 1b in Appendix A shows the median proportion of daily sales in each hour observed in our setting. (Going forward, unless stated otherwise we will use the terms "demand" and "sales" interchangeably to represent the sales revenue from core menu items before accounting for any discounts or taxes.)

While the actual demand forecasts made by the managers were not centrally tracked, preventing us from using them directly, we were given access to every restaurant's staffing roster in 2018. By combining these staffing data with the observed sales data in our sample and augmenting this with other covariates (e.g., seasonal factors, dates of promotional campaigns), we were able to reconstruct the managers' hourly demand forecasts by using a range of statistical forecasting methods with varying levels of sophistication. This is a commonly used technique when historic demand forecasts are unobserved and must be estimated from observed data (e.g., Rumyantsev and Netessine 2007, Freeman et al. 2020).

The fitted values from the statistical models become our estimate of *expected demand* for each hour of the day at each QSR in our sample. The more that actual demand (i.e., sales) deviates from expected demand (i.e., forecasted sales), the more likely it is that the QSR manager made incorrect staffing allocation and ingredient preparation decisions for that hour in the day, impacting operational performance.

Table 1 documents the comprehensive range of variables used in estimating demand. Not only do all of these variables influence demand, but their effects may not be uniform across all restaurants. For example, some restaurants may cater more to evening crowds and others more to lunchtime crowds, and the effectiveness of promotional campaigns may vary. To account for heterogeneity in the effect of each control variable across restaurants, our models also include two-way interactions between restaurant ID and all the other variables discussed above. This demand model is consistent with the bottom-up nature of demand forecasts, which are made by individual restaurant managers in response to local, restaurant-specific trends.

Table 1 Features used in training the demand model.

Variable	Type	Description
Hour*	Categorical(24)	Hour of the day
Day of the week	Categorical(7)	Day of the week on which the restaurant-day falls
Week number	Categorical(5)	$\lceil date/7 \rceil$, where $\lceil\rceil$ represents the *ceiling* function
Month	Categorical(12)	Month on which the restaurant-day falls
Public holiday	Binary	Takes a value of 1 if the restaurant-day falls on a public holiday
Holiday period	Binary	Takes a value of 1 if the restaurant-day falls during the prominent holiday season
Time trend	Continuous	Takes a value of 1 for the first day within a restaurant and increments by 1 for each subsequent day
Promo campaign	Binary	Takes a value of 1 if the restaurant-day falls during an active promotional campaign
Staffing data*	Continuous	Time-series of the number of staff that were on duty for each restaurant-hour
Restaurant ID	Categorical(99)	Unique ID to represent each QSR in our setting

Notes: Some of these features are also used as control variables in our econometric specification (cf. Section 6.1). Variables marked with an asterisk (*) are used only in constructing the demand model and are not included as controls in the regression models.
If a variable is categorical, the number in $(\cdot)$ in the 'Type' column indicates the number of levels.

However, a linear model encompassing all the variables described above along with their two-way interactions with restaurant ID would contain over 6500 features. To avoid over-fitting, we use a lasso (Tibshirani 1996) model for variable selection and regularization. We use the `cv.glmnet` function from the `glmnetUtils` package in `R` for model training and selection. The objective of lasso regression is to solve

$$\min_{\beta} \left\{ \sum_{i=1}^{N} \left(y_i - x_i^T \beta \right)^2 \right\} \quad \text{subject to} \quad \sum_{j=1}^{p} |\beta_j| \leq t \tag{1}$$

where y_i is each observed instance of hourly demand, x_i is the covariate vector for the i^{th} observation, and $\beta = (\beta_1, \ldots, \beta_p)$ is the set of coefficients to be estimated in the model. The regularization parameter for the lasso regression is given by t and is a free parameter that determines the amount of regularization.

In our models, we select the regularization parameter through five-fold cross-validation such that the mean absolute error (MAE) of the model is minimized. This method shrinks the total number of active features in our linear model by one-third – from 6643 to 4381. When trained on a randomly selected partition containing 70% of the data and tested on the remaining 30% of the data, such a model achieves an out-of-sample R-square of 69.5%. When trained on the whole set of observations, the fitted values from the lasso model, denoted $Demand^F_{idh}$, give us our estimates of demand at each restaurant i on each day d and each hour h. (See Appendix G for the results of analyses with a range of alternative models for building demand forecasts. These results are entirely consistent with those reported in the main paper).

Theoretically, restaurant demand can also be measured by number of diners served. Since the number of diners served by the QSR was not directly available to us, we compute this metric for each check by counting the number of core menu items that were sold as part of that check. Core menu items include all items except for desserts, beverages and sides. For instance, a check showing one meal – one burger, one side of fries and one beverage – would be considered as serving one diner. An a-la-carte check with two burgers, one side of fries and two beverages would be considered as serving two diners. We repeated our analysis with $Diners_{id}$ as the definition of demand and find the results to be fully consistent with our primary definition based on $Sales_{id}$.

5.3. Forecast Accuracy

Having constructed our measure of expected demand, $Demand^F_{idh}$, we can define our first key dependent variable: forecast accuracy. We define hourly errors, δ_{idh}, as the discrepancy between the observed and forecasted demand for each hour h on each day d at each restaurant i:

$$\delta_{idh} = Demand_{idh} - Demand^F_{idh} \tag{2}$$

We aggregate from the hourly level to the daily level to arrive at our measure of forecast error for each restaurant-day by taking the root mean squared error (RMSE) of these hourly errors, that is:

$$ForecastErr_{id} = \sqrt{\frac{\sum_{h=1}^{m} \delta^2_{idh}}{m}} \tag{3}$$

where m is the number of hours for which restaurant i was operational on day d. RMSE is commonly used to measure forecast accuracy and it is regularly used in the management literature (e.g., Harrison and Klein 2007, Ahuja et al. 2020).

Mean absolute error (MAE), $\sum_{h=1}^{m} |\delta_{idh}|/m$, is an alternative technique to measure forecast error – see Appendix G for results of analysis with this alternative technique that are fully consistent with those reported in the main paper.

5.4. High-Margin Item Sales

Our second dependent variable, $HighMargin_{id}$, relates to the composition of sales, with a focus on determining how platform dependence influences the sales revenue generated from high-margin items. In our analysis, high-margin items are defined as drinks, desserts, and sides. For any given restaurant i on day d, $HighMargin_{id}$ represents the fraction of total sales revenue attributable to high-margin items. Conceptually, it captures the restaurant's reliance on these profit-boosting items in its daily operations.

$HighMargin_{id}$ can also be defined as the proportion of total bills (checks) at restaurant i on day d that include at least one high-margin item. Analysis with this alternative definition produces consistent results and are reported in Table 3.

5.5. Dependence on Platforms

Our primary independent variable of interest is the QSR's dependence on third-party food-delivery platforms. Measuring a QSR's dependence on platforms is akin to asking its manager: "If today were a typical day, what proportion of today's demand would you expect to come from food-delivery platforms?" It is important to note that the answer to this question should not be affected by unpredictable factors that impact daily demand at a restaurant. In order to generate our measure of dependence, we start by calculating the proportion of total sales that come through third-party food-delivery platforms at each restaurant i on each day d, denoting this $PlatformShare_{id} = PlatformDemand_{id}/Demand_{id}$, where $PlatformDemand_{id}$ and $Demand_{id}$ give the sales made through the platform and sales through all channels, respectively.

Note that $PlatformShare_{id}$ is not a good measure of platform dependence since it is likely to be highly correlated with the idiosyncratic and unobservable factors that affect forecast accuracy. For example, if the weather is particularly bad on a given day or if traffic congestion is worse than normal, then overall demand might drop (i.e., forecast accuracy will be worse) and simultaneously the share of demand coming through the platform may increase (as customers may opt for the convenience of delivery over take-away or dining out). This situation would lead to endogeneity issues and bias our estimate of the impact of platform demand on forecast accuracy. We address these endogeneity concerns by eliminating the noise so that we can tease out each restaurant's latent dependence on platforms, Dep_{id}.

To do this, we compute a moving average of $PlatformShare_{id}$ over a sufficiently large window, that is,

$$Dep_{id} = \frac{\sum_{d \in T_d} PlatformShare_{id}}{|T_d|}. \tag{4}$$

Here, T_d is the set of dates that lie within the specified window surrounding day d (i.e., $T_d = \{d-w,\ldots,d-1,d+1,\ldots,d+w\}$, with w giving the number of days before and after day d that defines the length of the window) and $|T_d| = 2w-1$ denotes the cardinality of set T_d. Importantly, the window defined in Eq. (4) excludes day d in order to remove the impact of any idiosyncratic factors that might simultaneously affect dependence and demand forecast accuracy on a particular day. It should be noted that the QSR chain in our study ran a number of promotional campaigns in the time window when our data was gathered. Since these campaigns ran for multiple days, they may have had a non-random impact on $PlatformShare_{id}$, while also potentially impacting the accuracy of our demand forecasts on day d. While we already control for promotional campaigns in the demand forecast model, we are also careful to account for the impact of such campaigns when computing the moving averages to tease out Dep_{id}. Full details of the moving average specification are in Appendix B.

In essence, Dep_{id} is a moving average of the daily share of demand from platforms that serves as our measure for the underlying third-party food-delivery platform dependence of a particular restaurant i on that day d. When the window used is sufficiently long, this measure is stable (i.e., there are no large swings in dependence from one day to the next) but also follows the underlying trend in dependence at a particular restaurant (e.g., if dependence on the platform increases over our time horizon, then so too will this measure).

In this paper we use a wide time window, $w = 14$, when calculating Dep_{id}. This is equivalent to using a one-month window around day d and is chosen to reduce the impact of any residual unobserved factors on the days surrounding day d, as these factors might simultaneously affect dependence and demand forecast accuracy. What residual effects do exist will be accounted for in our control structure (specifically with seasonal factors). Estimations using other time windows (e.g., $w = 10$ and $w = 21$) are given in Appendix E, with results consistent with those reported when $w = 14$.

5.6. Summary Statistics

Summary statistics and correlations between the key variables outlined in Sections 5.3–5.5 are given in Table 2. To preserve the anonymity of the QSR chain, all summary statistics are reported on standardized variables, that is, after subtracting the mean and dividing by the standard deviation. However, statistical analysis and the results reported in the paper are for variables on their original scales. Note that $Sales_{id}$, $Diners_{id}$, and $PlatformShare_{id}$ do not enter any of the statistical models, but are provided for reference in Table 2. Of the remaining variables, the descriptive statistics show that all variables appear to be well-behaved except for $ForecastErr_{id}$, for which the maximum takes a value 9.65σ above the mean. This suggests that this variable is right skewed or contains one or more outliers. To reduce the influence of values in this right tail, we therefore take the natural logarithm transformation of $ForecastErr_{id}$ prior to modeling. Histograms of $ForecastErr_{id}$ are provided in Appendix A.

Table 2 Descriptive statistics and correlations for key variables.

Panel A: Descriptive Statistics – Standardized Variables

					St.Dev		
	Mean	Median	Max	Min	Overall	Between	Within
$Sales_{id}$	0	-0.21	14.49	-1.68	1	0.76	0.61
$Diners_{id}$	0	-0.22	12.06	-1.64	1	0.74	0.64
$ForecastErr_{id}$	0	-0.21	9.65	-1.85	1	0.44	0.88
$HighMargin_{id}$	0	0.01	4.29	-5.06	1	0.60	0.78
$PlatformShare_{id}$	0	-0.17	5.43	-1.26	1	0.70	0.66
Dep_{id}	0	-0.08	3.67	-1.48	1	0.84	0.50

Panel B: Correlations – Standardized Variables

	$Sales_{id}$	$Diners_{id}$	$ForecastErr_{id}$	$HighMargin_{id}$	$PlatformShare_{id}$	Dep_{id}
$Sales_{id}$	1.000					
$Diners_{id}$	0.959^{***}	1.000				
$ForecastErr_{id}$	0.603^{***}	0.619^{***}	1.000			
$HighMargin_{id}$	-0.049^{***}	-0.094^{***}	-0.110^{***}	1.000		
$PlatformShare_{id}$	-0.035^{***}	0.030^{***}	-0.033^{***}	-0.387^{***}	1.000	
Dep_{id}	0.002	0.052^{***}	0.004	-0.257^{***}	0.821^{***}	1.000

Notes: *p<0.1; **p<0.05; ***p<0.01; $Sales_{id}$ and $Diners_{id}$ refer, respectively, to the total observed daily sales and number of diners at each restaurant-site. The *Within* column depicts standard deviation within each restaurant-site, whereas the *Between* column depicts standard deviation across restaurant-sites in those variables. The *Overall* column depicts standard deviation in those variables pooled over the full dataset.

It should also be noted that Table 2 reports only a small correlation between Dep_{id} (our key independent variable) and forecast errors. However, this is misleading as these correlations do not account for confounding variables such as seasonality, trend and unobserved site-specific effects.

6. Models and Results

In this section we describe the estimation approach used to test Hypotheses 1 and 2 and then present the results.

6.1. Econometric Specification

Our empirical strategy leverages the variation in platform dependence within (rather than *between*) each restaurant over the period of observation. We use this variation to estimate the impact of platform dependence on errors in forecasting demand and the proportion of sales revenue from high-margin items. More formally, we test Hypotheses 1 and 2 by estimating the following set of independent regression equations:

$$ln(ForecastErr_{id}) = \gamma_i + \gamma_1 Dep_{id} + \boldsymbol{\gamma}_2^{\intercal}\mathbf{X}_{id} + \epsilon_{id}^{\gamma} \tag{5}$$

$$logit(HighMargin_{id}) = \phi_i + \phi_1 Dep_{id} + \boldsymbol{\phi}_2^{\intercal}\mathbf{X}_{id} + \epsilon_{id}^{\phi} \tag{6}$$

In Equations (5)–(6), the primary effects of interest are given by γ_1 and ϕ_1, which correspond to the effect of an increase in platform dependence on a QSR's error in demand forecast and proportion of sales revenue from high-margin items, respectively. Meanwhile, γ_i and ϕ_i represent the unobserved site-specific effects (i.e., the site-specific fixed effects) that account for unobserved differences between restaurant-sites that may impact the relationship between dependence and the outcome variables. The error terms ϵ_{id}^{γ} and ϵ_{id}^{ϕ} are normally distributed with a conditional mean value of zero.

The vector containing the control variables in Equations (5)–(6) is given by $\mathbf{X}_{id}$. This vector includes the controls listed in Table 1 and is also expanded to include an additional covariate to adjust for the impact of the scale of demand. In particular, as daily demand increases, so too might the errors in our forecasts (i.e., δ_{idh} in Equation (2) may increase). Hence, the forecast error rates may be higher on days with higher expected demand or in restaurants with larger overall demand. To account for this, the scale of demand at the restaurant-day level is included as a control in the model as the natural logarithm of the hourly forecasts aggregated across each day, that is, $Scale_{id} = ln(\sum_h Demand_{idh}^F)$.

6.2. Results

Table 3 reports the estimates of coefficients from Equations 5 and 6 outlined in Section 6.1. Standard errors reported are heteroskedasticity and autocorrelation consistent, and clustered at the restaurant-level.

Table 3 Regression results for Hypotheses 1 and 2

	Forecast Error		High Margin Sales	
	(1)	(2)	(3)	(4)
Dependence	0.283^{**}	0.314^{**}	-1.108^{***}	-2.638^{***}
	(0.096)	(0.099)	(0.090)	(0.195)
Seasonality	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes
Observations	32,695	32,695	32,695	32,695

Note: ***p<0.001; **p<0.01; *p<0.05
Standard errors in parentheses below each coefficient.

The first section in Table 3, titled "Forecast Error" reports results for Hypotheses 1a and 1b. Model 1 reports results for the outcome variable, $ForecastErr_{id}$, based on $Sales_{id}$, as defined in Section 5.3. Our results indicate that an increasing dependence on platforms has a significant impact on the error in forecasting demand (denoted by γ_1 in Equation (5)). For every 10 percentage point (p.p.) increase in platform dependence, error in forecasting demand increases by 2.83% (p-value < 0.01). In Model 2, we report corroborating results for a replication with an alternate definition where we use *number of diners served* (denoted by $Diners_{id}$) as the demand metric.

The second section in Table 3 reports results for Hypotheses 2a and 2b. Model 3 reports results for the outcome variable, $HighMargin_{id}$, as defined in Section 5.4. Our results indicate that an increasing dependence on platforms has a significant impact on the proportion of sales revenue from high-margin items (denoted by ϕ_1 in Equation (6)). For every 10 p.p. increase in platform dependence, the proportion of sales revenue from high-margin items shrinks by $1 - e^{(-0.1108)}$, i.e, it shrinks by 10.5% (p-value < 0.01). In Model 4, we report corroborating results for a replication where we define $HighMargin_{id}$ as the proportion of total orders (checks) at restaurant i on day d that include at least one high-margin item.

Overall, these results are consistent with evidence in support of Hypotheses 1b and 2b, rather than 1a and 2a. They demonstrate that the net effect of increasing dependence on platform demand is poorer demand forecasting performance and a smaller proportion of sales revenue from high-margin items.

7. Moderating Effects of Restaurant-Level Customer Characteristics

Despite the negative impact of food-delivery platforms on forecasting performance and high-margin sales, it is not easy for restaurants to stay off these platforms. Understanding the circumstances under which this impact is diminished or intensified will enable restaurants to manage their presence on these platforms. Therefore, in this section we identify and hypothesize possible moderating factors in our setting before discussing how these moderators are measured and presenting our models and results.

7.1. Moderators: Theory and Hypothesis

In their seminal review of the literature on restaurant management drawing on five decades of scholarship, Thompson (2010) systematically identifies six key dimensions of customer demand. This influential and widely cited review is an authoritative source for understanding the intricacies of consumer behavior in restaurant settings. In addition to demand arrival patterns and items purchased, the paper lists the following key characteristics of customer demand at restaurants:

- *Party size*: Number of customers at each table. Some restaurants are more popular with larger groups whereas others may have a more intimate ambience.
- *Party composition*: This can encapsulate many dimensions such as the age, gender, affluence and ethnicity of the customers.
- *Meal duration*: The time elapsed between a party being seated at a table and the table being ready for the next party.
- *Peak distribution*: The distribution of demand through the restaurant's opening hours. For instance, some restaurants are busy during lunch and less crowded during dinner whereas others maybe equally busy during both lunch and dinner windows.

While our core investigation focuses on forecast accuracy and the sale of high-margin items, the inter-linked nature of these characteristics suggests that the above factors may moderate the impact of third-party food-delivery platforms on our dependent variables. Unfortunately, our dataset does not allow us to observe meal durations. However, we can provide a more nuanced understanding of the implications of platform dependence by studying the moderating effects of party size, party composition (modeled as affluence) and peak distribution (modeled as lunch vs. dinner demand ratio). We hypothesize that:

Hypothesis 3: *Customer characteristics, such as affluence, party size and business vs. leisure, moderate the relationship between increasing dependence on third-party food-delivery platforms and demand characteristics*

7.2. Moderators: Variable Construction

Affluence We operationalize the concept of "affluence" at the restaurant level through the average sales revenue obtained per core menu item per check. This provides an insight into the spending capacity of customers frequenting each restaurant.

Party Size "Party Size" is identified at the restaurant level by calculating the mean number of core menu items ordered per check. This metric provides an indirect estimate of the size of customer groups visiting the restaurant, with the underlying assumption that larger groups tend to order more core menu items.

Peak Distribution To understand the temporal dynamics of customer visits, we define the "Peak Distribution" at each restaurant as the ratio of checks served during the typical lunch hours (11am to 2pm) to those during the dinner window (7pm to 10pm). To ensure the ratio is symmetric, we apply the natural logarithm to the ratio. The resulting variable takes a value of zero when lunch and dinner footfalls are identical. A negative value indicates a predominant dinner crowd, whereas a positive value signals a busier lunch period.

For ease of interpretation, we report results with moderators as standardized variables, that is, after subtracting the mean and dividing by the standard deviation. Histograms of these moderating variables are provided in Appendix A. There are alternative ways to operationalize our moderating variables. For instance, *Affluence* and *Party Size* can be defined as the median sales revenue obtained from core menu items and the median number of core menu items per check, respectively. Further, *Peak Distribution* can be defined as an indicator variable for restaurants in which more checks are served at lunch than at dinner. Analyses with these alternative definitions, which produces nearly identical results, are available in Appendix C.

7.3. Moderators: Econometric Specification and Results

Model Specification: To test Hypothesis 3, which focuses on the moderating effect of customer demand characteristics, we include additional interaction terms in the regression equations:

$$ln(ForecastErr_{id}) = \gamma_i + \gamma_1 Dep_{id} + \boldsymbol{\gamma}_2^\intercal \mathbf{X}_{id} + \boldsymbol{\gamma}_3^\intercal Dep_{id} \times \boldsymbol{\lambda}_i + \epsilon_{id}^{\gamma} \tag{7}$$

$$logit(HighMargin_{id}) = \phi_i + \phi_1 Dep_{id} + \boldsymbol{\phi}_2^\intercal \mathbf{X}_{id} + \boldsymbol{\phi}_3^\intercal Dep_{id} \times \boldsymbol{\lambda}_i + \epsilon_{id}^{\phi} \tag{8}$$

In Equations (7)–(8), $\boldsymbol{\lambda}_i$ is a vector containing the three site-level characteristic features of customer-attributes that we hypothesize moderate the relationship between platform dependence and the two outcomes. The vector of coefficients $\boldsymbol{\gamma}_3$ and $\boldsymbol{\phi}_3$ indicates the extent of the moderating effects. We include these moderators one at a time, resulting in three separate regression equations being estimated for each dependent variable.

Results: Table 4 reports the estimates of coefficients from Equations 7 and 8. Standard errors reported are heteroskedasticity and autocorrelation consistent, and clustered at the restaurant-level.

The models titled "High-margin Sales" represent the results of estimating the regression model Equation 8. The coefficients for *Affluence*, *Party Size* and *Peak Distribution* can be interpreted as the moderating effect of a one standard deviation increase in each customer characteristic on the relationship between platform dependence and high-margin sales. The positive estimate (0.180; p<0.001) for the coefficient of *Affluence* means that it mitigates the decrease in high-margin sales associated with increasing dependence. In other words, for the same amount of dependence, a restaurant frequented by a more affluent customer base will experience a smaller decrease in the proportion of sales revenue from high-margin items. Specifically, a restaurant with affluence two standard deviations above the mean will see a 7.2% ($= 1 - e^{-0.1106+2\times0.0180}$) reduction in high-margin sales for every

Table 4 Regression results with Moderating Factors

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.283**	0.283**	0.272**	0.246*	−1.108***	−1.106***	−1.166***	−1.159***
	(0.096)	(0.096)	(0.088)	(0.101)	(0.090)	(0.083)	(0.090)	(0.092)
× Affluence	–	−0.015	–	–	–	0.180***	–	–
		(0.064)				(0.036)		
× Party Size	–	–	0.027	–	–	–	0.148*	–
			(0.091)				(0.059)	
× Peak Distribution	–	–	–	−0.092†	–	–	–	−0.130*
				(0.053)				(0.054)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,695	32,695	32,695	32,695	32,695	32,695	32,695	32,695
Adjusted/Pseudo R^2	0.348	0.348	0.348	0.348	0.286	0.286	0.286	0.286

Note: ***p<0.001; **p<0.01; *p<0.05; †p<0.1
Standard errors in parentheses below each coefficient.

10 percentage point increase in dependence, compared to a 13.6% reduction for a restaurant two standard deviations below the mean. This demonstrates that customer traits significantly moderate the impact of dependence on high-margin item sales. Notably, even the most affluent restaurant in our dataset, at 3.76 standard deviations above the mean, experiences decreased high-margin sales as dependence increases.

Similarly, the positive estimate (0.148; $p<0.05$) for the coefficient of *Party Size* means that it mitigates the decrease in high-margin sales associated with increasing dependence. In other words, for the same amount of dependence, a restaurant that typically sees larger parties will experience a smaller decrease in the the proportion of sales revenue from high-margin items. Lastly, the estimate for the coefficient of *Peak Distribution* is negative (-0.130; $p<0.05$). This means that restaurants that are more popular during the lunch peak than at dinner will experience a larger decrease in the proportion of sales revenue from high-margin items. As with affluence, even the restaurant that typically caters to the largest party sizes (2.79 standard deviations above the mean) and the restaurant that has the lowest peak distribution (2.25 standard deviations below the mean) experience a decrease in high-margin sales as dependence increases.

In Table 4, the models titled "Forecast Error" represent results for estimating the regression model Equation 7. We find no evidence to support the hypothesis that *Affluence* or *Party Size* moderates the relationship between platform dependence and forecast error. However, we find weak evidence to suggest that restaurants that are more popular during the lunch peak than at dinner will experience a smaller increase in forecast error. The estimate for the coefficient of *Peak Distribution* is negative (-0.092; $p<0.1$) but we do not find this moderating effect to be robust against alternative model specifications. These results suggest that the adverse impact of platform dependence on demand forecast accuracy is largely independent of customer characteristics.

8. Alternative Explanations and Robustness Tests

It is seldom straightforward to establish causality through the analysis of secondary data. Nevertheless, we try to eliminate several other possible explanations for the effects we observe in our data using a variety of approaches described below. We also perform a wide variety of robustness checks to ensure that our results and insights are not confined to the specifications presented in the main manuscript.

First, by taking advantage of only the within-restaurant variation in dependence to estimate our models, we control for a wide variety of unobserved site-specific effects, such as restaurant size, neighborhood demographics, and the intensity of competitor activity. Additionally, there were no technology or ownership changes during the period of our study. There could have been some changes to competitor activity at a few restaurants but since these changes are not *systematic* in nature, they would not materially affect our estimates.

Second, our findings are based on the underlying degree of *dependence* each restaurant has on demand that comes through food-delivery platforms. To eliminate the impact of such confounding factors as sudden changes in weather or traffic or local events such as sporting events, the calculation for this measure does not include day d, and the measure is calculated over a wide time window to further reduce potential confounding by factors that persist for multiple days. We have also repeated our analysis employing a range of window sizes to measure dependence on platforms (Appendix E). To further account for confounding by such multi-day factors, we also (i) perform robustness checks by excluding additional days around the focal day when calculating dependence (Appendix F) and (ii) estimate a model that includes the lag of the dependent variable on the right-hand side in addition to those features presented in Table 1 (Appendix D). The latter technique is frequently used to correct biases in coefficient estimates due to omitting a potential time-varying confounding variable (see e.g. Gokpinar et al. 2010).

Third, as econometricians, we are not privy to on-ground information about the demand at a particular QSR beyond that which is reported in our data set. However, we use staffing data as a proxy for on-ground information about the demand at a particular QSR and augment it with other demand features, allowing us to reconstruct forecasts that closely reflect the managers' hourly demand forecats. Furthermore, we build a wide range of demand models with varying levels of sophistication, to ensure that our findings are robust to various forecasting models (Appendix G).

Finally, to verify the robustness of our findings against variations in variable operationalization, we replicate our analysis with alternative measures. These have been described throughout the paper, and include, e.g., Mean Absolute Error (MAE) for forecast error (Appendix G), the proportion of bills containing at least one high-margin item as an alternative for High-Margin sales (Table 3), and different metrics for the moderating variables (Appendix C).

Overall, regardless of the demand model specified or the way we measure dependence, forecast error, and high-margin-item sales, we consistently find that as the QSR increases its dependence on third-party platforms, its error in forecasting demand increases and its proportion of revenue from high-margin items decreases. We also find robust evidence that the impact of platform dependence on high-margin sales is moderated by affluence, group size, and peak distribution.

9. Conclusion

Using a novel, proprietary transaction-level dataset this paper investigates how an increasing dependence on third-party food-delivery platforms impacts two key demand characteristics for restaurants: forecast accuracy and high-margin item sales. Our analysis reveals that greater platform dependence is associated with an adverse impact on both these characteristics.

Our analysis shows that a 10 percentage-point increase in the proportion of a restaurant's sales from food-delivery platforms is associated with a 2.83% increase in demand forecast error. Accurate demand forecasts are critical for maintaining operational efficiency, and in the restaurant business, operational efficiency is a key driver of profitability (Roy et al. 2022). Therefore, this reduced ability to accurately predict demand can have major implications for restaurant managers. Overestimating demand will lead to excess food procurement and labor costs, while underestimating demand results in lost sales and poor customer experience.

Additionally, our results indicate that sales revenue from high-margin items, such as drinks, desserts, and sides, shrinks by 10.5% when platform dependence rises by 10 p.p. These items are key drivers of profit margins, so this shift in order composition is potentially problematic. Managers should therefore re-optimize menu design, pricing, and promotional strategies to shore up high-margin item sales through delivery platforms. They may also consider partnerships with specialized vendors to provide complementary profitable items suited to delivery.

Our analysis also reveals how restaurant-level characteristics moderate the relationship between platform dependence and high-margin sales. We find the negative impact is weaker for restaurants with higher customer spend or larger order sizes, indicating more affluent patrons and larger parties help preserve profitable item sales. However, restaurants with a greater demand during the lunch peak hours than at dinner suffer more severe effects. This suggests platform dependence may disproportionately affect restaurants that rely on daytime sales from small, not-so-affluent groups such as retail employees or students.

These insights equip restaurant managers with a more nuanced understanding of the impact of platform dependence on their performance. With this knowledge, they may be better positioned to navigate negotiations with food-delivery platforms. This might entail advocating for lower commission rates for restaurants with certain customer characteristics or forging strategic alliances that could buffer against the potential adverse outcomes of heightened platform dependence.

Importantly, our results suggest that there are costs for restaurants associated with food-delivery platforms, beyond the commission charges. For instance, food-delivery platforms have occasionally listed menus without a restaurant's consent (Bellan 2022). In these cases, even if the restaurants do not pay commissions on the take-away orders generated from the platforms, since the same kitchen must also satisfy the platform-based orders, these restaurants would experience a deterioration in their forecast accuracy and profit margins.

Additionally, these findings indicate that the cloud kitchen business model may not be as profitable as previously anticipated. Even though these businesses may save a considerable amount of overhead costs, our insights suggest that their demand would be much more volatile and that they will likely sell a smaller proportion of high-margin items than a traditional restaurant. There is anecdotal evidence to this effect in popular media (Singh 2020).

Our study also marks a crucial step in broadening the academic dialogue around platform dependence and its impact on demand characteristics in the restaurant industry. We complement the ongoing conversation on commissions and cannibalization by offering empirical evidence of the complex ways in which platform dependence can influence other key demand characteristics. Our work emphasizes the malleability of these demand characteristics and suggests that future research should consider these dynamics when modelling the intricate interplay between businesses and digital platforms.

In conclusion, understanding the unintended adverse consequence of food-delivery platforms is crucial for the overall effective managements of restaurants. Although the rise of food-delivery platforms ushers in a unique set of challenges, these obstacles can be effectively managed through a proactive approach to understanding and strategically addressing the ramifications of increased platform dependence. Our study lays a robust empirical foundation for such strategic decision-making, offering a significant contribution to the burgeoning discourse on digital platforms in the field of operations management. We encourage future research to further refine and expand upon our findings, enhancing our collective understanding of platform dependence, its drivers, and its wide-ranging implications.

References

Ahuja K, Chandra V, Lord V, Peens C (2021) Ordering in: The rapid evolution of food delivery. *McKinsey & Company* URL `https://www.mckinsey.com/business-functions/sustainability/our-insights/sustainability-blog/how-ai-can-unlock-a-127b-opportunity-by-reducing-food-waste`.

Ahuja V, Alvarez CA, Staats BR (2020) Maintaining continuity in service: An empirical examination of primary care physicians. *Manufacturing & Service Operations Management* 22(5):1088–1106.

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 Operations Management* 13(4):489–507.

Andrews RL, Currim IS (2004) Behavioural differences between consumers attracted to shopping online versus traditional supermarkets: implications for enterprise design and marketing strategy. *International Journal of Internet Marketing and Advertising* 1(1):38–61.

Astashkina E, Belavina E, Marinesi S (2019) The environmental impact of the advent of online grocery retailing. *Available at SSRN 3358664* .

Bakos JY (1991) A strategic analysis of electronic marketplaces. *MIS Quarterly* 15(3):295.

Bakos JY (1997) Reducing buyer search costs: Implications for electronic marketplaces. *Management science* 43(12):1676–1692.

Bakos Y (1998) The emerging role of electronic marketplaces on the internet. *Communications of the ACM* 41(8):35–42.

Banerjee S, Johari R, Riquelme C (2016) Dynamic pricing in ridesharing platforms. *ACM SIGecom Exchanges* 15(1):65–70.

Bapna R, Chang SA, Goes P, Gupta A (2009) Overlapping online auctions: Empirical characterization of bidder strategies and auction prices. *MIS quarterly* 763–783.

Bapna R, Goes P, Gupta A, Jin Y (2004) User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS quarterly* 21–43.

Belavina E, Girotra K, Kabra A (2017) Online grocery retail: Revenue models and environmental impact. *Management Science* 63(6):1781–1799.

Bell DR, Gallino S, Moreno A (2018) Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Science* 64(4):1629–1651.

Bellan R (2022) Uber eats to pay millions for listing chicago restaurants without consent. *Techcrunch* URL `https://techcrunch.com/2022/12/05/uber-eats-to-pay-millions-for-listing-chicago-restaurants-without-consent/`.

Brynjolfsson E, Smith MD (2000) Frictionless commerce? a comparison of internet and conventional retailers. *Management science* 46(4):563–585.

Burt S, Sparks L (2003) E-commerce and the retail process: a review. *Journal of Retailing and Consumer Services* 10(5):275–286.

Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management* 19(3):368–384.

Cachon GP, Terwiesch C (2016) *Operations Management* (McGraw-Hill).

Chandrashekar TS, Narahari Y, Rosa CH, Kulkarni DM, Tew JD, Dayama P (2007) Auction-based mechanisms for electronic procurement. *IEEE Transactions on Automation Science and Engineering* 4(3):297–321.

Chang HH, Sokol DD (2022) How incumbents respond to competition from innovative disruptors in the sharing economy—the impact of airbnb on hotel performance. *Strategic Management Journal* 43(3):425–446.

Chen M, Hu M, Wang J (2022) Food delivery service and restaurant: Friend or foe? *Management Science* 68(9):6539–6551.

Chen RR, Roundy RO, Zhang RQ, Janakiraman G (2005) Efficient auction mechanisms for supply chain procurement. *Management Science* 51(3):467–482.

Chen T, Guestrin C (2016) Xgboost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.

Chen Y, Hu M, Zhou Y (2019) Pricing and matching in the sharing economy. *Sharing Economy*, 137–164 (Springer).

Cramer J, Krueger AB (2016) Disruptive change in the taxi business: The case of uber. *American Economic Review* 106(5):177–82.

Cui H, Niu R, Tong X (2022) Food in the time of covid: Estimating the effects of delivery platforms on small-business restaurants in california. *Available at SSRN 4219542* .

De Vries J, Roy D, De Koster R (2018) Worth the wait? how restaurant waiting time influences customer behavior and revenue. *Journal of operations Management* 63:59–78.

Dewey C (2017) Analysis — the insane $43 billion system that gets food delivered to your door — washingtonpost.com. URL `https://www.washingtonpost.com/news/wonk/wp/2017/08/08/the-insane-43-billion-system-that-gets-food-delivered-to-your-door/`.

Dunn E (2018) How delivery apps may put your favorite restaurant out of business. *The New Yorker* URL `https://www.newyorker.com/culture/annals-of-gastronomy/are-delivery-apps-killing-restaurants`.

Feldman P, Frazelle AE, Swinney R (2023) Managing relationships between restaurants and food delivery platforms: Conflict, contracts, and coordination. *Management Science* 69(2):812–823.

Forman L (2019) Some restaurants get served in delivery boom. *The Wall Street Journal* .

Freeman M, Robinson S, Scholtes S (2020) Gatekeeping, fast and slow: An empirical study of referral errors in the emergency department. *Management Science* .

Friedman JH (2001) Greedy function approximation: a gradient boosting machine. *Annals of statistics* 1189–1232.

Geurts P, Ernst D, Wehenkel L (2006) Extremely randomized trees. *Machine learning* 63(1):3–42.

Gokpinar B, Hopp WJ, Iravani SM (2010) The impact of misalignment of organizational structure and product architecture on quality in complex product development. *Management science* 56(3):468–484.

H2Odocs (2020) Automl: Automatic machine learning. *H2O.ai documentation* URL `https://docs.h2o.ai/h2o/latest-stable/h2o-docs/automl.html`.

Haddon H, Jargon J (2019) Consumers love food delivery. restaurants and grocers hate it. *The Wall Street Journal* .

Harrison DA, Klein KJ (2007) What's the difference? diversity constructs as separation, variety, or disparity in organizations. *Academy of management review* 32(4):1199–1228.

Houck B (2017) Why some restaurants are cutting ties with mobile ordering apps.

Isaac M, Yaffe-Bellany D (2019) The rise of the virtual restaurant. *The New York Times* URL `https://www.nytimes.com/2019/08/14/technology/uber-eats-ghost-kitchens.html`.

Kapoor R, Agarwal S (2017) Sustaining superior performance in business ecosystems: Evidence from application software developers in the ios and android smartphone ecosystems. *Organization Science* 28(3):531–551.

Kumar A, Mehra A, Kumar S (2019) Why do stores drive online sales? evidence of underlying mechanisms from a multichannel retailer. *Information Systems Research* 30(1):319–338.

Li Z, Wang G (2020) The role of on-demand delivery platforms in restaurants. *Available at SSRN 3813891* .

Liaw A, Wiener M, et al. (2002) Classification and regression by randomforest. *R news* 2(3):18–22.

Lunden I (2020) Restaurant rewards booking app seated nabs $30m, acquires venuebook to add events. *Techcrunch* URL `https://techcrunch.com/2020/08/18/restaurant-rewards-booking-app-seated-nabs-30m-acquires-venuebook-to-add-events/`.

Lynch Jr JG, Ariely D (2000) Wine online: Search costs affect competition on price, quality, and distribution. *Marketing science* 19(1):83–103.

Lyu G, Cheung WC, Teo CP, Wang H (2019) Multi-objective online ride-matching. *Available at SSRN 3356823* .

Magnin C (2019) How ai can unlock a usd 127b opportunity by reducing food waste. *McKinsey & Company* URL `https://www.mckinsey.com/business-functions/sustainability/our-insights/sustainability-blog/how-ai-can-unlock-a-127b-opportunity-by-reducing-food-waste`.

Mahajan A (2019) How to survive the competitive online food delivery business. *Hackernoon Magazine* URL `https://hackernoon.com/how-to-survive-in-competitive-online-food-delivery-business-b26a6ce8a0e0`.

Malone TW, Yates J, Benjamin RI (1987) Electronic markets and electronic hierarchies. *Communications of the ACM* 30(6):484–497.

Mithas S, Jones JL (2007) Do auction parameters affect buyer surplus in e-auctions for procurement? *Production and Operations Management* 16(4):455–470.

Nie YM (2017) How can the taxi industry survive the tide of ridesourcing? evidence from shenzhen, china. *Transportation Research Part C: Emerging Technologies* 79:242–256.

Özkan E, Ward AR (2020) Dynamic matching for real-time ride sharing. *Stochastic Systems* 10(1):29–70.

Parboteeah DV, Valacich JS, Wells JD (2009) The influence of website characteristics on a consumer's urge to buy impulsively. *Information systems research* 20(1):60–78.

Pereira DXR (2018) Going zero waste in canteens: Exploring food demand using data analytics .

Ransom D (2009) Can they really make money off the dollar menu? *Wall Street Journal* .

Roy D, Spiliotopoulou E, de Vries J (2022) Restaurant analytics: Emerging practice and research opportunities. *Production and Operations Management* 31(10):3687–3709.

Rumyantsev S, Netessine S (2007) What can be learned from classical inventory models? a cross-industry exploratory investigation. *Manufacturing & Service Operations Management* 9(4):409–429.

Ryu S, Kim K, Hahn J (2022) Crowdfunding success effects on financing outcomes for startups: A signaling theory perspective. *MIS Quarterly (Accepted)* .

Samuel J, Zheng ZE, Xie Y (2020) Value of local showrooms to online competitors. *MIS Quarterly* 44(3).

Schlosser E (2012) *Fast food nation: The dark side of the all-American meal* (Houghton Mifflin Harcourt).

Sharma A, Mehrotra A (2007) Choosing an optimal channel mix in multichannel environments. *Industrial Marketing Management* 36(1):21–28.

Shufflebotham B (2022) Mcdonald's behind-the-scenes secrets shared by kitchen insider. *QSR Magazine* URL `https://www.walesonline.co.uk/whats-on/food-drink-news/went-kitchen-mcdonalds-learnt-behind-24630406`.

Siddharth Cavale HR (2021) Crunch time: Fast-food chains look to keep food crisp as deliveries soar — reuters.com. URL `https://www.reuters.com/article/us-fast-food-delivery-crispy/crunch-time-fast-food-chains-look-to-keep-food-crisp-as-deliveries-soar-idUSKBN2BN1CX`.

Simchi-Levi D, Kaminsky P, Simchi-Levi E (2003) *Designing and managing the supply chain: Concepts, strategies, and case studies* (McGraw-Hill).

Singh M (2020) India's swiggy to cut 1,100 jobs, scale down cloud kitchen operations. *Techcrunch* URL `https://techcrunch.com/2020/05/18/indias-swiggy-to-cut-1100-jobs-scale-down-cloud-kitchen-operations/x`.

Smith MD, Bailey J, Brynjolfsson E (1999) *Understanding digital markets: review and assessment* (MIT press).

Soh C, Markus ML, Goh KH (2006) Electronic marketplaces and price transparency: strategy, information technology, and success. *MIS quarterly* 705–723.

Standing C, Love PE, Stockdale R, Gengatharen D (2006) Examining the relationship between electronic marketplace strategy and structure. *IEEE Transactions on Engineering Management* 53(2):297–311.

Standing S, Standing C, Love PE (2010) A review of research on e-marketplaces 1997–2008. *Decision Support Systems* 49(1):41–51.

Swaminathan JM, Lee HL (2003) Design for postponement. *Handbooks in Operations Research and Management Science* 11:199–226.

Tan TF, Netessine S (2020) At your service on the table: Impact of tabletop technology on restaurant performance. *Management Science* 66(10):4496–4515.

Thompson GM (2010) Restaurant profitability management: The evolution of restaurant revenue management. *Cornell Hospitality Quarterly* 51(3):308–322.

Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1):267–288.

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

Wang S, Zheng S, Xu L, Li D, Meng H (2008) A literature review of electronic marketplace research: Themes, theories and an integrative framework. *Information Systems Frontiers* 10(5):555–571.

Wiggers K (2018) Tenzo raises usd1.8 million to reduce restaurant food waste with ai. *VentureBeat* URL `https://venturebeat.com/2018/11/05/tenzo-raises-1-8-million-to-reduce-restaurant-food-waste-with-ai/`.

Xia Y, Zhang GP (2010) The impact of the online channel on retailers' performances: An empirical evaluation. *Decision Sciences* 41(3):517–546.

Zervas G, Proserpio D, Byers JW (2017) The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Journal of marketing research* 54(5):687–705.

Appendix
A. Plots and Histograms

In this section, we provide some plots and histograms of the key variables used in our paper.

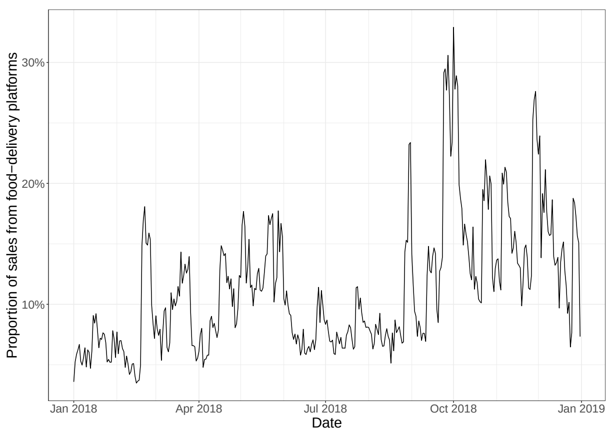

(a) Median $PlatformShare_{id}$ over time

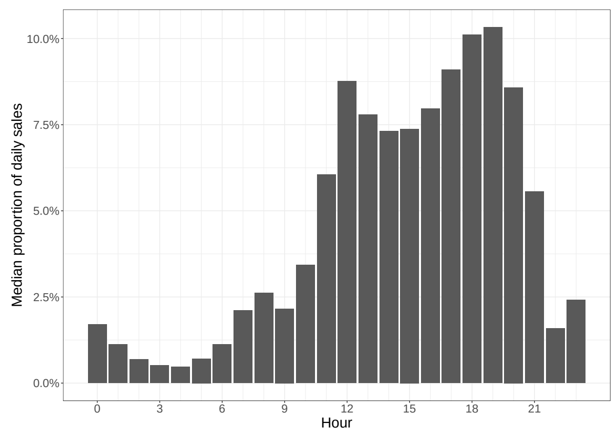

(b) Median proportion of daily sales in each hour

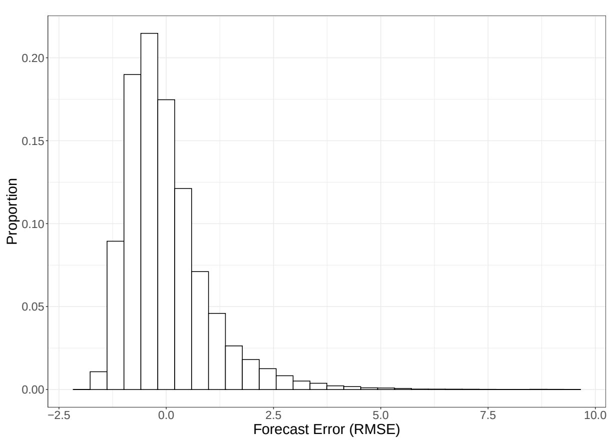

(c) Histogram of $ForecastErr_{id}$

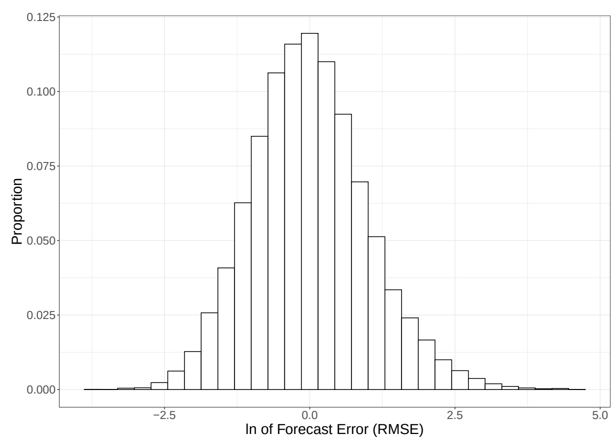

(d) Histogram of $ln(ForecastErr_{id})$

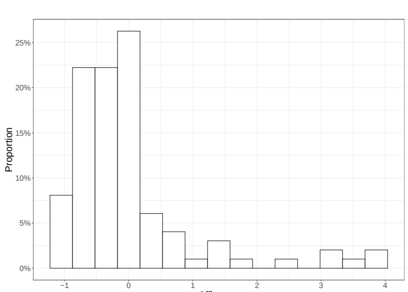

(e) Histogram of Affluence

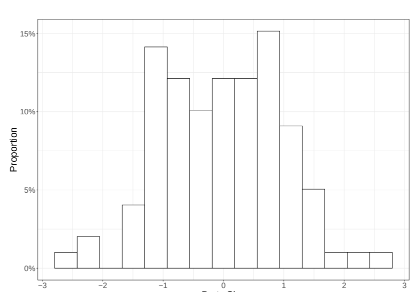

(f) Histogram of Party Size

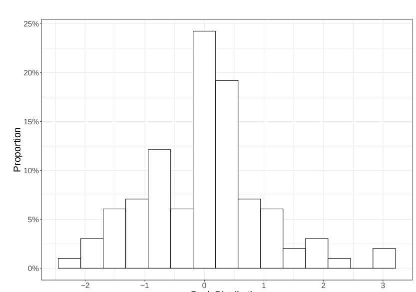

(g) Histogram of Peak Distribution

Figure 1 Plots and histograms of key variables

B. Measuring Dependence on Platforms

We define the share of demand coming from third-party food-delivery platforms for each restaurant-day as

$$PlatformShare_{id} = \frac{PlatformDemand_{id}}{Demand_{id}}$$

As we noted in Section 5.5, $PlatformShare_{id}$ is not a good measure of platform *dependence* since it is likely to be impacted by idiosyncratic and unobservable factors. For example, if the weather or traffic is particularly bad on a given day, the share of demand coming through the platform may increase (as customers may opt for the convenience of delivery rather than having to collect their meal themselves), i.e.,

$$PlatformShare_{id} = Dep_{id} + noise_{id} \tag{9}$$

We tease out the latent variable, Dep_{id}, by eliminating the noise around it. We compute Dep_{id} as a moving average of $PlatformShare_{id}$ over a sufficiently large window. However, note that the QSR chain ran a number of promotional campaigns in the time window to which our data belongs. Since these campaigns ran for multiple days, they may have had a non-random impact on $PlatformShare_{id}$, and a simple moving average of $PlatformShare_{id}$ would not lead us to Dep_{id}. We therefore carefully account for the impact of promotional campaigns when computing the moving averages to tease out Dep_{id}. We do this by estimating the following regression equation for every day d, at every restaurant i, in our sample:

$$PlatformShare_{it} = \lambda_{id} + \gamma_{id} \times banner.binary_{it} + \Delta_{it} \tag{10}$$

where $t \in [d-w, d+w] \wedge t \neq d$ and w is the size of the window.

We use λ_{it} as our measure of dependence for restaurant i on day d. Note that if we do not control for the banner activities, then the intercept λ_{it} is equivalent to a sample average of $PlatformShare_{it}$ over that window. In the main paper, we present results with $w = 14$ days (cf. Table 3 and Table 4). In Appendix E and Appendix F we present results with alternate window definitions.

C. Alternative measures of moderating variables

In Section 7.2 we describe the operationalization of our moderating variables. There are alternative ways to operationalize our moderating variables. In this section, we report analyses with these alternative definitions, which produces nearly identical results as reported in Sections 6.2 and 7.3.

As an alternative measure, we define *Affluence* and *Party Size* as the median sales revenue obtained from core menu items and the median number of core menu items per check, respectively. We define *Peak Distribution* as an indicator variable for restaurants in which more checks are served at lunch than at dinner.

Table 5 Results with alternative definitions for moderators

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.283**	0.282**	0.277**	0.262*	−1.108***	−1.097***	−1.138***	−1.146***
	(0.096)	(0.096)	(0.094)	(0.099)	(0.090)	(0.083)	(0.088)	(0.089)
× Affluence	–	−0.023	–	–	–	0.223***	–	–
		(0.072)				(0.046)		
× Party Size	–	–	0.020	–	–	–	0.099†	–
			(0.064)				(0.052)	
× Peak Distribution	–	–	–	−0.054	–	–	–	−0.100*
				(0.057)				(0.048)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,695	32,695	32,695	32,695	32,695	32,695	32,695	32,695
Adjusted/Pseudo R^2	0.348	0.348	0.348	0.348	0.286	0.286	0.286	0.286

Note: ***p<0.001; **p<0.01; *p<0.05; †p<0.1
Standard errors in parentheses below each coefficient.

D. Regression with Lagged Dependent Variable

In this section, we present results from regressions that include the first lag of forecast error as an additional covariate, along with those presented in Table 1).

Table 6 Results from models that include first lag of DV a control

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.224^{*}	0.224^{*}	0.211^{**}	0.194^{*}	-0.560^{***}	-0.560^{***}	-0.585^{***}	-0.585^{***}
	(0.088)	(0.088)	(0.080)	(0.095)	(0.042)	(0.039)	(0.042)	(0.043)
× Affluence	–	−0.028	–	–	–	0.072^{***}	–	–
		(0.050)				(0.018)		
× Party Size	–	–	0.033	–	–	–	0.062^{*}	–
			(0.082)				(0.029)	
× Peak Distribution	–	–	–	−0.077	–	–	–	-0.063^{*}
				(0.048)				(0.026)
First lag of DV	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,170	32,170	32,170	32,170	32,695	32,695	32,695	32,695
Adjusted/Pseudo R^2	0.353	0.353	0.353	0.353	0.286	0.285	0.285	0.285

Note: $^{***}p<0.001$; $^{**}p<0.01$; $^{*}p<0.05$; $^{\dagger}p<0.1$
Standard errors in parentheses below each coefficient.

E. Alternative Window Sizes for Computing Dependence

In our main results, we chose a ±14-day window to run the regressions. In this section, we show that our results are robust to other window size specifications. Table 7 showcases results for regressions using a ±10-day window, which is approximately equivalent to averaging over 1.5 weeks in both directions: past and future.

Table 8 showcases results from a re-run of our analysis, this time using a ±20 day moving window to measure Dep_{id}. This is approximately equivalent to averaging over three weeks in both directions: past and future.

Table 7 Results with Dependence defined over a three weeks moving window

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.228**	0.227**	0.213**	0.193*	−1.041***	−1.034***	−1.093***	−1.085***
	(0.084)	(0.085)	(0.078)	(0.088)	(0.078)	(0.071)	(0.077)	(0.079)
× Affluence	–	−0.038	–	–	–	0.187***	–	–
		(0.063)				(0.033)		
× Party Size	–	–	0.040	–	–	–	0.148**	–
			(0.087)				(0.054)	
× Peak Distribution	–	–	–	−0.097†	–	–	–	−0.124*
				(0.054)				(0.052)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,874	32,874	32,874	32,874	32,874	32,874	32,874	32,874
Adjusted/Pseudo R^2	0.347	0.347	0.347	0.347	0.286	0.286	0.286	0.286

Note: ***p<0.001; **p<0.01; *p<0.05; †p<0.1
Standard errors in parentheses below each coefficient.

Table 8 Results with Dependence defined over a six weeks moving window

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.288*	0.289*	0.282**	0.244†	−1.112***	−1.113***	−1.174***	−1.170***
	(0.119)	(0.119)	(0.106)	(0.128)	(0.102)	(0.096)	(0.103)	(0.105)
× Affluence	–	−0.044	–	–	–	0.181***	–	–
		(0.065)				(0.037)		
× Party Size	–	–	0.013	–	–	–	0.147*	–
			(0.091)				(0.063)	
× Peak Distribution	–	–	–	−0.101†	–	–	–	−0.133*
				(0.058)				(0.055)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,296	32,296	32,296	32,296	32,296	32,296	32,296	32,296
Adjusted/Pseudo R^2	0.349	0.349	0.349	0.349	0.286	0.286	0.286	0.286

Note: ***p<0.001; **p<0.01; *p<0.05; †p<0.1
Standard errors in parentheses below each coefficient.

F. Alternative Definitions of Dependence

In Appendix B we discuss the moving average specification employed to measure Dep_{id}. Specifically, we note that the moving average is defined over a window $t \in [d-w, d+w] \wedge t \neq d$ and $w = 14$ for the results presented in the main paper (cf. Table 3 and Table 4).

In the above definition, we disregard only the focal day d to eliminate the impact of confounding factors. To test the robustness of our results, we compute variants of *dependence* that also disregard $PlatformShare$ from a few days surrounding the focal day d.

More formally, we estimate the regression specified in Eq. 10 over $t \in [d-w, d+w] \wedge t \notin [d-l, d+l]$, where $w = 14$ and $l \in \{1, 2, 3\}$.

Table 9 Results with an alternative definition of Dependence ($w = 14$ and $l = 1$)

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.298**	0.298**	0.286**	0.261*	−1.076***	−1.073***	−1.131***	−1.124***
	(0.098)	(0.099)	(0.089)	(0.104)	(0.088)	(0.082)	(0.088)	(0.090)
× Affluence	–	−0.023	–	–	–	0.171***	–	–
		(0.066)				(0.036)		
× Party Size	–	–	0.030	–	–	–	0.143*	–
			(0.090)				(0.058)	
× Peak Distribution	–	–	–	−0.094†	–	–	–	−0.124*
				(0.054)				(0.053)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,486	32,486	32,486	32,486	32,486	32,486	32,486	32,486
Adjusted/Pseudo R^2	0.349	0.349	0.349	0.349	0.286	0.286	0.286	0.286

Note: Standard errors in parentheses below each coefficient. ***p<0.001; **p<0.01; *p<0.05; †p<0.1

Table 10 Results with an alternative definition of Dependence ($w = 14$ and $l = 2$)

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.321**	0.321**	0.309***	0.289**	−1.042***	−1.039***	−1.094***	−1.088***
	(0.096)	(0.096)	(0.088)	(0.102)	(0.085)	(0.079)	(0.085)	(0.087)
× Affluence	–	−0.020	–	–	–	0.165***	–	–
		(0.066)				(0.035)		
× Party Size	–	–	0.031	–	–	–	0.136*	–
			(0.091)				(0.056)	
× Peak Distribution	–	–	–	−0.084	–	–	–	−0.120*
				(0.053)				(0.051)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,488	32,488	32,488	32,488	32,488	32,488	32,488	32,488
Adjusted/Pseudo R^2	0.349	0.349	0.349	0.349	0.286	0.286	0.286	0.286

Note: Standard errors in parentheses below each coefficient. ***p<0.001; **p<0.01; *p<0.05; †p<0.1

Table 11 Results with an alternative definition of Dependence ($w = 14$ and $l = 3$)

	Forecast Error				High-margin Sales			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Dependence	0.338***	0.338**	0.324***	0.311**	−1.005***	−1.000***	−1.052***	−1.048***
	(0.097)	(0.097)	(0.091)	(0.104)	(0.083)	(0.077)	(0.083)	(0.084)
× Affluence	–	−0.025	–	–	–	0.161***	–	–
		(0.067)				(0.034)		
× Party Size	–	–	0.035	–	–	–	0.128*	–
			(0.092)				(0.055)	
× Peak Distribution	–	–	–	−0.072	–	–	–	−0.114*
				(0.054)				(0.050)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,278	32,278	32,278	32,278	32,278	32,278	32,278	32,278
Adjusted/Pseudo R^2	0.349	0.349	0.349	0.350	0.286	0.286	0.286	0.286

Note: Standard errors in parentheses below each coefficient. ***p<0.001; **p<0.01; *p<0.05; †p<0.1

G. Alternative Measures of Forecast Error

This section presents results from alternate demand models, using two forecast error measures: RMSE and MAE (Table 12). The "Main Model" columns use the demand model from Section 5.2, with the "RMSE" column same as reported in Table 3.

The "Labor Only Model" columns in Table 12 are based on a simpler model including only staffing data and restaurant-specific factors, excluding variables in Table 1. It simulates less sophisticated QSR managers' intuition-based demand estimation, yielding an out-of-sample R-square of 42.6%.

Finally, the "Stacked Model" columns in Table 12 are based on a more sophisticated model using machine learning and ensemble methods to combine the forecasts of several algorithms. We use `h2o.automl` function from the `h2o` package in `R` to facilitate model training and selection, allowing for automatic tuning and training of machine learning models with minimal user input (H2O.docs 2020). The best model, chosen through five-fold cross validation such that the mean absolute error (MAE) of the model is minimized, was a stacked ensemble of models such as gradient boosting machines (Friedman 2001), XGBoost models (Chen and Guestrin 2016), fully-connected multi-layer artificial neural networks, random forests (Liaw et al. 2002), and extremely randomized trees (Geurts et al. 2006). Each of these constituent models was automatically tuned by `h2o` to minimize MAE. This is clearly a more advanced method than is used by restaurant managers in our QSR sites, and it is able to increase out-of-sample R-square to 78.8%.

Regardless of the demand model specified or the way we measure forecast error we consistently find that as the QSR increases its dependence on third-party platforms, its error in forecasting demand increases.

Table 12 Results from models that use alternative measures of Forecast Error

	Main Model		Labor Only Model		Stacked Model	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
Dependence	0.283^{**}	0.273^{**}	0.251^{*}	0.227^{*}	0.326^{***}	0.340^{***}
	(0.096)	(0.099)	(0.132)	(0.135)	(0.089)	(0.081)
Seasonality	Yes	Yes	Yes	Yes	Yes	Yes
Holiday Period	Yes	Yes	Yes	Yes	Yes	Yes
Day Aggregate Forecast	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend	Yes	Yes	Yes	Yes	Yes	Yes
Promo Campaign	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,695	32,695	33,302	33,302	33,308	33,308

Note: $^{***}p<0.001$; $^{**}p<0.01$; $^{*}p<0.05$
Standard errors in parentheses below each coefficient.