The Business School for the World®

Working Paper
2020/35/TOM

An Unintended Consequence of Platform Dependence: Empirical Evidence from Food-Delivery Platforms

Varun Karamshetty
INSEAD, varun.karamshetty@insead.edu

Michael Freeman
INSEAD, michael.freeman@insead.edu

Sameer Hasija
INSEAD, sameer.hasija@insead.edu

Last revised date was 27 July 2020

Food waste is a severe economic and social problem. Restaurants contribute significantly to food waste because they face the classic trade-off between speed of service and leftover inventory, which is particularly crucial in the context of quick service restaurants (QSRs). To offer a high speed of service, QSRs pre-cook most of their food, but they can hold it only for a short time. To effectively manage this trade-off, QSRs have become increasingly reliant on demand forecasts. However, online food-delivery platforms that connect restaurants, riders/drivers, and consumers are growing in popularity, and it is unclear how the growth of food-delivery platforms impacts the ability of restaurants to accurately forecast their demand. We empirically investigate the impact of food-delivery platforms on the demand forecast error in QSRs and analyze the underlying mechanism. We find that as customers become increasingly dependent on food-delivery platforms, QSR demand becomes harder to forecast. We also find that the majority of the increase in overall forecast error is due to an increase in the error associated with the demand pattern and a smaller portion is due to error in forecasting demand amplitude. Based on our results, we offer suggestions for QSRs on how to manage their relationship with food-delivery platforms to decrease their forecast error and increase operational effciency.

Keywords: Food-Delivery Platforms; Quick Service Restaurants; Forecasting; Econometrics

An Unintended Consequence of Platform Dependence: Empirical Evidence from Food-Delivery Platforms

Varun Karamshetty
INSEAD, Singapore, varun.karamshetty@insead.edu

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

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

Food waste is a severe economic and social problem. Restaurants contribute significantly to food waste because they face the classic trade-off between speed of service and leftover inventory, which is particularly crucial in the context of quick service restaurants (QSRs). To offer a high speed of service, QSRs pre-cook most of their food, but they can hold it only for a short time. To effectively manage this trade-off, QSRs have become increasingly reliant on demand forecasts. However, online food-delivery platforms that connect restaurants, riders/drivers, and consumers are growing in popularity, and it is unclear how the growth of food-delivery platforms impacts the ability of restaurants to accurately forecast their demand. We empirically investigate the impact of food-delivery platforms on the demand forecast error in QSRs and analyze the underlying mechanism. We find that as customers become increasingly dependent on food-delivery platforms, QSR demand becomes harder to forecast. We also find that the majority of the increase in overall forecast error is due to an increase in the error associated with the demand *pattern* and a smaller portion is due to error in forecasting demand *amplitude*. Based on our results, we offer suggestions for QSRs on how to manage their relationship with food-delivery platforms to decrease their forecast error and increase operational efficiency.

Key words: Food-Delivery Platforms; Quick Service Restaurants; Forecasting

---

The sad fact is that 30–40 percent of all food is wasted each year, yet as many as one in eight people suffer from chronic hunger.
—Nicolas Burquier, Chief Customer and Operations Officer, Pizza Hut

1. Introduction

Food waste is an urgent economic and social problem on a global scale. The Sustainable Development Goals 2030 charter, adopted by the United Nations in 2015, explicitly includes food waste in its objectives, calling for efforts to "halve per capita global food waste at the retail and consumer level, and reduce food losses along production and supply chains by 2030." Each year, food losses and waste amount to US$1 trillion – more than 1% of the global GDP. The hospitality and food service sector is a significant contributor to this problem, and restaurants are the second largest contributor to food waste in the United States (cf. ReFED 2016, Gunders 2017).

Food waste is a result of the classic trade-off that restaurants make between too much and too little. This problem is particularly egregious for quick service restaurants (QSRs) – an industry with estimated revenues nearing US$300 billion in 2019[1] – as QSRs pre-cook most of their food and rely on speed of service as a key competitive advantage. Although operational issues in the restaurant industry have previously been studied (cf. De Vries et al. 2018, Pereira 2018, Ülkü et al. 2019), the rise of smartphone-enabled food-delivery platforms is beginning to change the food service industry landscape and therefore compels a fresh look at this topic.

Consumers who are already used to ordering books, clothing and accessories online are now embracing the idea of ordering their food online as well. Advances in digital technology (such as ubiquitous smartphones, API connectivity, and low-cost mobile data) have facilitated the growth of online food-delivery platforms that generate revenue by connecting restaurants, riders/drivers, and consumers. Globally, the estimated food-delivery market reached a value of US$94.4 billion in 2019, and it is expected to grow at a rate of 9.3% CAGR to reach US$134.5 billion by 2023 (Klein 2019). Demand from delivery platforms is often the fastest growing demand channel for QSRs (Khan 2019).

The growth of the food-delivery market has been further accelerated by the lockdown measures implemented in response to the COVID-19 pandemic (Chiappetta 2020), as food-delivery platforms have proven essential in cities around the globe. Even though restaurants' dependence on these platforms is likely to scale back slightly post-lockdown, it is unlikely to go back to the pre-lockdown levels for two reasons. First, patrons are hesitant to visit restaurants and might continue to order home-delivered food (Awasthi 2020, First 2020). Second, forced exposure may have accustomed consumers to the advantages of home-delivery platforms (e.g., convenience, choice) over dine-in services, thus permanently affecting consumer behavior. This creates an urgent necessity to understand the operational implications of consumer participation in these platforms.

[1] Revenues of the US fast food industry since 2002. (Statista)

While it is well-understood that platforms potentially increase a firm's market size and reach to its customers (Sharma and Mehrotra 2007, Xia and Zhang 2010), it is unclear how they affect the firm's operational performance. In particular, it is not clear how these platforms impact two important operational measures for QSRs: customer waiting times and food waste. Moreover, many QSRs do not have systems that reliably measure these performance metrics on a continuous basis, preventing a rigorous study of these second-order effects of food delivery platforms on restaurants. However, an operations management (OM) lens can be a convenient tool to study these new business models even without such detailed operational data: Fundamental OM principles dictate that the accuracy of demand forecasting at a QSR will be a crucial factor that affects both food waste and customer service levels. Measures of forecast accuracy can be constructed from the more reliable (and easily available) sales data, clearing the path to study operational implications of food delivery platforms on QSRs.

Given the criticality of demand forecasts in a QSR's ability to balance food waste and responsiveness, it becomes important to understand how increasing dependence on such platforms may impact demand forecast accuracy. Yet there are competing hypotheses about how food-delivery platforms can affect a QSR's ability to forecast its demand. On the one hand, delivery platforms reduce the impact of day-to-day noise on demand. When customers order food from the comfort of their homes without regard for the conditions outside (e.g., weather, traffic), the restaurant's demand may become more stable and easier to predict. On the other hand, delivery platforms may also increase heterogeneity among customers' order times (for instance, by removing the socialization aspect of dining in), thus making the restaurant's demand more noisy and harder to predict.

Additionally, on the one hand, delivery platforms increase the catchment area of a restaurant. This should pool the restaurant's risk over a larger population of customers and decrease the relative error in forecasting demand on any given day. On the other hand, a customer browsing the delivery platform's mobile application (or website) for dining options can effortlessly switch between restaurants. This increases the level of competition amongst restaurants (Mahajan 2019) and may increase the relative error in forecasting demand for any given day. In the end, it is not clear how delivery platforms impact the accuracy of a restaurant's demand forecasts, which in turn will impact the restaurant's overall operational performance.

The aim of this paper is to understand how and why food-delivery platforms impact a QSR's ability to forecast demand. We use a proprietary database of detailed transaction-level data from a QSR chain. The database has a record of every transaction made in each restaurant in the chain in 2018, totaling approximately 50 million transactions. For each transaction, the database also has information on the mode of delivery to the customer: dine-in, take-away, drive-through or delivery platform.

Our results show that a 10 percentage point increase in a QSR's dependence on delivery platforms leads to a 2.83% increase in its overall forecast error. This has significant implications in terms of food waste, speed of service and operational efficiency for QSRs that are already finding it difficult to justify the commissions paid out to delivery platforms (Dunn 2018, Cagle 2020, Tkacik 2020). These findings could be particularly significant for the many restaurants that were forced to go purely online during the COVID-19 outbreak: a simple linear extrapolation of our results suggests that some restaurants may have seen a nearly one-third increase in forecast inaccuracy, leading to either substantial increase in food waste, decline in service performance, or both.

Further, we show that the majority of this increase in overall forecast error is due to an increase in the error of forecasted demand *pattern*, i.e., the distribution of demand within a day. A smaller portion of the phenomenon is due to an increase in error of the forecasted demand *amplitude*, i.e., volume on a specific day.[2] Overall, our study suggests that in order to mitigate the negative effect of delivery platforms on forecast accuracy, restaurants should focus their efforts on stabilizing their demand patterns.

From an academic research perspective, our results have implications for future work related to the operations of food-delivery platforms, as the decrease in forecast accuracy should be accounted for in studies that model the relationship between restaurants and online food-delivery platforms. From a practical perspective, our analysis of the underlying mechanism allows us to offer suggestions for restaurants to attenuate the adverse effect of food-delivery platform dependence on forecastability.

2. Contributions to the Literature

There is growing interest in understanding the strategies that underpin business models for the app-based sharing economy and e-marketplaces. This is particularly the case for OM

[2] We provide a detailed explanation of the terms *amplitude* and *pattern* in Section 3.

researchers, who have studied various aspects: the fundamental theory of e-marketplaces, optimization of underlying systems and technology, paths to adoption by participants, and operational implications of participation.

The fundamental theory of e-marketplaces (and electronic commerce, more generally) received significant attention even before the recent advances in digital technology made e-marketplaces ubiquitous. Early work by Malone et al. (1987) predicted that advances in technology would shift more economic activity into e-marketplaces. In subsequent years, a robust body of literature developed to model how electronic marketplaces can reduce transaction costs (Bakos 1997), resulting in reduced intermediation costs and more efficient inter-organizational transactions (Bakos 1991), ultimately propelling us towards "frictionless" markets (Bakos 1998). These papers were followed by empirical studies that tested the theoretical hypotheses and added nuance to our understanding of the theory of e-marketplaces (e.g. Smith et al. 1999, Brynjolfsson and Smith 2000).[3]

More recently, as economic activity in different sectors did indeed shift to e-marketplaces (e.g., ride sharing, food delivery, gig workers), researchers responded by studying ways to optimize the underlying systems of a marketplace. This rich stream of literature looks at interconnected aspects of e-marketplace system design, such as auction types and strategies (e.g., Bapna et al. 2004, Mithas and Jones 2007, Bapna et al. 2009), procurement (e.g., Chen et al. 2005, Chandrashekar et al. 2007), pricing (e.g., Banerjee et al. 2016, Cachon et al. 2017) and order dispatching (e.g., Chen, Hu and Zhou 2019, Lyu et al. 2019, Özkan and Ward 2020). Yet while there is a preponderance of such research into system design, there are few articles studying the implications for e-marketplace participants (Kapoor and Agarwal 2017).

In fact, in their review of literature on e-marketplaces, Standing et al. (2010) find that this discrepancy constitutes a research gap: there is too much emphasis on the e-marketplace and its workings and not enough on the impact of such business models on participating firms. There are a few studies evaluating these business models from the participant's strategic perspective (e.g., Soh et al. 2006) and examining the benefits and costs associated with participation (e.g., Standing et al. 2006). However, the operational

[3] A comprehensive review of the literature on the theory of e-marketplaces is outside the scope of this paper, and we refer the reader to Wang et al. (2008) and Standing et al. (2010).

implications of e-marketplaces on participating firms still remain largely understudied, and our paper contributes to this much-needed conversation.

Rather than examining the inner workings of app-based platforms (e-marketplaces), we focus on how engaging with them is impacting established businesses, and we do this in the particular context of third-party food-delivery platforms and QSRs. In doing so, we build on the work of Feldman et al. (2018), who were the first to study the relationship between food-delivery platforms and traditional restaurants. Using a queuing model with a single observable stream of customers who choose between dining at the restaurant and ordering on the food-delivery platform, they find that the commonly observed one-way revenue contract yields inefficient outcomes from a system perspective.

Chen, Hu and Wang (2019) follow up by studying an unobservable queuing model with two streams of customers that are heterogeneous in their tech-savviness. They find that food-delivery platforms do not necessarily increase the overall demand for the restaurant, but they might cannibalize demand from traditional channels as the segment of tech-savvy customers grows. Our study complements these papers by empirically establishing a different adverse consequence (reduced forecastability of demand) of a restaurant's dependence on third-party food-delivery platforms.

Our paper also complements the literature on multichannel retail businesses. When a restaurant chooses to participate in an e-marketplace, it opens a new sales channel, thus transforming into a multichannel sales enterprise. In a non-restaurant context, such multichannel systems have been extensively studied to maximize the efficiency of inventory procurement, warehousing and distribution to satisfy demand in both channels (see Burt and Sparks (2003) for a review of relevant literature).

For multichannel sales enterprises, the ubiquity of low-cost, high-speed mobile data and high-performance processors is blurring the line between offline and online demand, as customers are now able to switch between these channels seamlessly. This has prompted researchers to study how processes in one channel can impact processes in another – for instance, Gallino et al. (2017) and Akturk et al. (2018) study multichannel operations for retailers with a "buy-online, ship-to-store" system, and Kumar et al. (2019) study how the presence of an offline store impacts online sales for an apparel retailer. We refer the reader to Agatz et al. (2008) and Zhang et al. (2010) for reviews of literature associated with this theme.

Although the literature of multichannel retailing that examines the interrelationships between channels is relevant to our study, the restaurant industry differs from other retail in an important way: the finished goods (cooked food items, in our context) are highly perishable. This means that demand from the online channel (food-delivery platform) cannot be pooled across stores and satisfied via drop-shipping or through a "buy-online, ship-to-store" policy. While pooling demand across time and geographies might be possible for online grocery retailers (cf. Belavina et al. 2017, Astashkina et al. 2019), the short shelf life of finished goods (and the impatience of hungry customers) forces restaurants to schedule all production at the store level.

Since restaurants regularly deal with finished goods that have a short shelf life, they try hard to optimize their operations to minimize waste while maintaining a good quality of service. Especially for a QSR, where speed of service is a key indicator of quality of service and margins are razor thin, operational efficiency plays an important role in achieving competitive advantage. In fact, improving efficiency through standardization of operational procedures has been crucial in the in the worldwide success of QSR chains like McDonald's

It is thus important for a QSR to accurately forecast its demand. Demand forecasts can help regulate customers' waiting time, which in turn affects consumption (Ülkü et al. 2019) and restaurant revenues (De Vries et al. 2018). Accurate demand forecasts also play an important role in reducing a restaurant's food waste (Pereira 2018). However, very little work to date directly studies the factors that impact demand forecasts, and we fill this gap by studying how third-party food-delivery platforms are impacting traditional restaurants' ability to accurately forecast their demand. To the best of our knowledge, we are the first to study this problem – empirically or otherwise.

3. Hypothesis Development

Since restaurants are in the business of selling highly perishable goods to impatient customers, they must navigate a common trade-off inherent to all businesses dealing with demand uncertainty. In this context, that means finding the right balance between speed of service and food waste. Pre-preparing too much food improves responsiveness but increases food waste. On the other hand, not pre-preparing enough food reduces food waste but hurts speed of service. In this scenario, accurate and granular demand forecasts can help a restaurant achieve the right balance (Wiggers 2018, Magnin 2019).

The granularity with which these forecasts must be made varies from restaurant to restaurant and depends largely on a customer's willingness to wait. In restaurants offering sit-down service with a waiter, customers are comfortable with (and sometimes even expect) a wait time of few minutes between placing their order and being served. These restaurants typically optimize their menu such that several dishes have recipes that are largely similar except for the final few steps. This allows staff to pre-prepare the base for several dishes in bulk and then execute the final steps to differentiate into individual dishes only after an order is placed (commonly known as *postponement* in the OM literature; cf. Swaminathan and Lee (2003)). Typically, the pre-prepared base has a longer holding time than the finished dish and is prepared only once or twice per day – removing the need for hourly (or more frequent) forecasts.

In contrast, customers visiting QSRs are disinclined to wait. For instance, Allon et al. (2011) find that customers at QSRs value their wait time at several times the average wage in the US. QSR chains are cognizant of this and routinely compete to offer shorter wait times (Pittman 2019). They achieve this by standardizing their operational procedures and pre-preparing most of their food even before an order comes in. Since such pre-prepared food has a very short holding time (typically only a few minutes), it is even more imperative that QSRs are able to accurately identify peaks and troughs in demand throughout the day to ensure that ingredients are ready when demand materializes. An accurate demand forecast will be able to identify this *pattern* and thus help the restaurant better match supply (e.g., staff, ingredient preparation) and demand.

In addition to accurately predicting the pattern of demand within a day, it is also important that restaurants are able to accurately estimate the total demand that will materialize on a particular day. If daily demand is anticipated to be higher than its actual value, this can result in various inefficiencies including overstaffing, unnecessary preparation of product components, and squandered utilities (Cachon and Terwiesch 2016, p. 390-391). On the other hand, underestimating demand can result in high workload in the kitchen, delays in service (Thompson 1998), and an increase in avoidable errors (Tan and Netessine 2014). Therefore, restaurants must also be able to accurately identify the scale of demand, i.e., its *amplitude*, when forecasting.

Given the importance of both amplitude and pattern when constructing a demand forecast within a QSR, the next sections discuss the anticipated impact of a food-delivery platform on each of these two metrics.

3.1. Demand Amplitude

By participating in a third-party food-delivery platform, a traditional restaurant is establishing an online sales channel in parallel to its existing offline channel, giving customers the option to seamlessly switch between channels according to their personal preferences. The advantage of opening an online sales channel is that it increases a firm's reach to its customers (Sharma and Mehrotra 2007), effectively increasing its market size (Xia and Zhang 2010). As the pool of potential customers increases, it can lead to a lower relative error in forecasting demand amplitude – a phenomenon extensively studied in the risk pooling literature (cf. Levi et al. 2003).

On the flip side, food-delivery platforms adversely impact the competitive landscape in which a restaurant operates. Any restaurant can potentially participate on a platform, including so-called *virtual kitchens* that do not have any offline presence (cf. Isaac and Yaffe-Bellany 2019). In addition, instead of only competing with other restaurants in its geographic area, a restaurant opening an online channel is now competing with others that have the same delivery area, which may be different and also larger. Thus a restaurant may gain many competitors that do not exist or do not compete with them in the offline space. Moreover, delivery platforms' apps and websites are designed to effortlessly compare prices, customer ratings and menu choices, allowing customers to easily switch between restaurants while browsing – which can increase price sensitivity among consumers (Lynch Jr and Ariely 2000). An increase in the number of competing restaurants, together with increased price sensitivity and very low switching costs for the customer, can make the restaurant's day-to-day demand more volatile and harder to predict.

On balance, then, it is not clear how third party food-delivery platforms will impact a restaurant's ability to accurately forecast demand *amplitude*, so we pose it as an empirical question to be answered in this paper:

Empirical Question 1: *As dependence on third-party food-delivery platforms increases, how does it impact the forecast error in a QSR's demand amplitude?*

3.2. Demand Pattern

QSRs plan their within-day operations (staffing, inventory to pre-prepare, etc.) based on their demand pattern forecasts. However, demand patterns can be disrupted by several exogenous, unforeseeable factors. For example, sudden changes in weather or traffic conditions can deter a customer from physically traveling to a restaurant. But customers in the restaurant's online sales channel are not impacted by such exogenous factors. In fact, there may be a negative correlation between demand coming from the online and physical channels, attenuating the variability in demand. As such, as a restaurant depends increasingly on delivery platforms, patterns in its overall demand should become easier to forecast.

However, online channels may also expose restaurants to more heterogeneity in customers' order times. Offline, customers' restaurant visits may be channeled into certain predictable times of day because customers are socializing with others (Andaleeb and Conway 2006) or because of workplace restrictions. But delivery platforms nullify such factors and remove the time constraints, giving customers the flexibility to order and receive food based on their personal preferences. This may make it difficult to accurately forecast the pattern of demand within a day and may increase the error in forecasting demand patterns.

It is therefore not clear how third party food-delivery platforms will impact a restaurant's ability to accurately forecast their demand *pattern*, so we ask:

Empirical Question 2: *As dependence on third-party food-delivery platforms increases, how does it impact the forecast error in a QSR's demand pattern?*

3.3. Total Forecast Error

It is possible for a QSR to accurately forecast demand *amplitude* but misallocate demand *pattern*. In this case, it will experience food shortages (and consequently low speed of service) during the periods where demand was underestimated and food waste during the periods where demand was overestimated. A QSR that accurately forecasts the day's demand pattern but overestimates the demand amplitude will waste food in every period on that day. On the other hand, getting the demand pattern right but underestimating demand amplitude will lead to longer waiting times and lower speed of service throughout the day.

If food-delivery platforms impact errors in both demand pattern and demand amplitude so that they move in the same direction (e.g., they both increase), then overall forecast error should also be impacted in that direction. However, if error in demand pattern and error in demand amplitude do not move in the same direction, it is not clear how overall forecast error will change. Even if both pattern and amplitude forecast error move in the same direction, they may have different relative strengths in influencing the overall forecast error. We therefore pose two more empirical questions:

Empirical Question 3: *How does increasing dependence on third-party food-delivery platforms impact the error in a QSR's overall demand forecast?*

Empirical Question 4: *What is the relative importance of demand pattern and demand amplitude in mediating the impact of third-party food-delivery platforms on the error in a QSR's overall demand forecast?*

4. Data Description and Variable Construction

4.1. Empirical Setting

We answer the empirical questions outlined in Section 3 by taking advantage of a proprietary Point of Sale (PoS) database. The database was obtained from a chain of QSRs, and it consists of detailed transaction-level data for approximately 50 million orders placed in 99 QSRs spread across the country in 2018. [4] 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. According to industry experts, this platform 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, there was another smaller platform that 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 will use these orders to generate our measure of each QSR's dependence on third-party food-delivery platforms (cf. Section 4.4). We refer to the other 88.9% of the orders as *non-platform demand*. Non-platform demand is distributed as follows: 51.1% dine-in, 27.8% take-away, and 6.1% drive-through (numbers reported as 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.

[4] Location is undisclosed to preserve anonymity.

4.2. Forecasting Demand

In our setting 486,211 hourly forecasts were made across all restaurants in the QSR chain over the one-year sample period. 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 almost exclusively use *revenue generated through sales before discounts* (denoted by $Sales_{id}$) as the metric for measuring demand. This makes sense because in this context, the sales price for each SKU is highly correlated with 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. Consistent with this, we use sales as the primary metric to measure demand in this paper.[5] (Going forward, unless stated otherwise we will use the terms "demand" and "sales" interchangeably to represent the value of sales generated 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 this 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 can recreate the managers' hourly demand forecasts by using a range of statistical forecasting methods of varying levels of sophistication. This is a commonly used technique in the literature when historic demand forecasts are unobserved and must be estimated from observed data (e.g., Rumyantsev and Netessine 2007, Freeman et al. 2017).

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 performance in the ways outlined in Section 1.

[5] However, we also replicate our analysis using *number of diners served* (denoted by $Diners_{id}$) as the demand metric – results can be found in Appendix C.1 and are consistent with those in the main paper.

In the rest of this section we document the comprehensive range of variables used in estimating demand (see Section 4.2.1) and describe the model specification that we use as our baseline approach in the paper (see Section 4.2.2).

4.2.1. Control Variables Here we document and justify the various covariates (summarized in Table 1) used to forecast demand.

Seasonality and trend. As one might expect, demand at any restaurant will vary across different hours of the day, with demand typically high during traditional mealtimes (e.g., lunch, dinner) and lower at other times of the day. Demand can also vary by a wide margin depending on the day of the week – for instance, a restaurant located in a shopping mall will typically experience higher demand on weekends. There is also variation in demand across weeks in the month, with demand tending to pick up during weeks that contain a payday. Additionally, at different times of the year, demand may vary due to factors such as the amount of daylight and peaks in tourism/travel.

To account for this variation, we include the following variables in our forecasting models: (i) hour of the day, (ii) day of the week, (iii) week number within that month[6] and (iv) month of the year. We also include a linear trend variable at the day level in our forecasting models to account for any systematic trends in demand over time.

Holiday period. To account for a significant anticipated shift in demand that occurs during a prominent holiday period, we include a binary variable that takes value one while that holiday is occurring and zero otherwise. We also include a separate binary variable that takes a value one during public holidays.

Promotional campaigns. In 2018, the QSR chain ran several nationwide promotional campaigns aimed at both online and offline customers. While we do not have detailed information regarding these campaigns, we were provided with the dates on which they were active. We include this information in our forecasting model as a binary input that takes a value of one when there is an active promotional campaign on that restaurant-day and zero otherwise. (Including a separate binary for each campaign does not improve out of sample accuracy of our model.)

[6] This is measured as $\lceil date/7 \rceil$, where $\lceil . \rceil$ represents the *ceiling* function

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 5.1). 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.

Staffing data. As noted before, QSR managers in our setting staff their restaurants based on their hourly sales forecasts. They use a mix of permanent and temporary employees in four- or eight-hour shifts to manage their hourly staffing levels. We use this staffing data (i.e., the number of staff scheduled to work during each hour of each day at each restaurant) as a proxy for on-ground information about the demand at a particular QSR, which the manager is privy to but we the econometricians are not. Such on-ground information might include, for instance, information about conferences or sporting events in the vicinity, public holidays, local issues such as elections, and various other factors that might affect demand at a QSR.

Restaurant-specific factors. Finally, a wide range of restaurant-level factors could affect demand. These include location, competitive landscape, and neighborhood demographics. To account for variation in demand across restaurants, we include the restaurant ID as a categorical variable in our model.

Not only do all of these control variables influence demand, but their effects may not be uniform across all restaurants. For example, some restaurants may cater more to evening crowds, others more to lunchtime crowds. Or promotional campaigns may be more effective for some restaurants than others. 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 other variables discussed above. This is consistent with the bottom-up nature of demand forecasts, which are made by individual restaurant managers in response to local, restaurant-specific trends.

4.2.2. Demand Model Specification Here we outline the baseline empirical model that we use to estimate hourly demand at each of the QSR sites in our dataset.

First, we note that 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 thus 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 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%.

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.[7] In this paper we will report results where demand is estimated using the lasso approach described here. However, in the appendix we also explore a range of alternative models for building demand forecasts. These models include:

1. A simpler model including only staffing data and restaurant-specific factors as features and ignoring the other variables listed in Table 1. This model provides an out-of-sample R-square of 42.6%, and it captures a scenario where QSR managers are less sophisticated and rely primarily on intuition and experience to estimate demand rather than using statistical models. Analysis with this model is described in Appendix C.2.

2. A more sophisticated model using machine learning and ensemble methods to combine the forecasts of several algorithms (e.g., gradient boosting and random forests) into a single aggregate forecast. 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%. A full discussion of this approach and replication of the analysis are described in Appendix C.3.

[7] We use the superscript F to denoted fitted or forecasted values as opposed to observed values.

Results using these alternative approaches are entirely consistent with those reported in the rest of this paper.

4.3. Measures of Forecast Accuracy

Now that we have constructed our measure of expected demand, $Demand^F_{idh}$, we are ready to define three key variables for our analysis: measures of forecast accuracy. Recall that we are interested in three primary measures: (1) *overall* forecast accuracy on a particular day, which we then break down into (2) accuracy in forecasting the demand *amplitude* on a day and (2) accuracy in forecasting the demand *pattern* within a day. In the rest of this section we outline the calculation of these variables, with histograms provided in Figure 1.

4.3.1. Overall Forecast Error We define hourly errors, δ_{idh}, as the discrepancy between 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 overall measure of forecast error for each restaurant-day by taking the root mean squared error (RMSE) of these hourly errors, i.e.:

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

where m is the number of hours for which restaurant i was operational on day d.[8] RMSE is a common metric used to measure forecast accuracy and it is regularly used in the management literature (e.g., Harrison and Klein 2007, Ahuja et al. 2019).

4.3.2. Relative Error in Demand Amplitude Amplitude error captures the extent to which we under- or over-predict total demand on a particular day when aggregating over our hourly demand forecasts. This is equal to the relative deviation of observed demand amplitude from forecasted demand amplitude, i.e.,

$$AmpErr_{id} = \frac{|Amp^F_{id} - Amp_{id}|}{Amp^F_{id}}. \tag{4}$$

In this equation, observed demand amplitude on a given day is simply the sum of observed hourly demand on that day, i.e., $Amp_{id} = \sum_{h \in h_{id}} Demand_{idh}$; forecasted demand amplitude

[8] Mean absolute error (MAE), $\sum_{h=1}^{m} |\delta_{idh}|/m$, is an alternative technique to aggregate hourly errors to measure overall forecast error – see Appendix D for analysis with this alternative technique, which produces nearly identical results.

is the sum of forecasted hourly demand on that day, i.e., $Amp^F_{id} = \sum_{h \in h_{id}} Demand^F_{idh}$; and h_{id} is the set of hours h over which restaurant i was open on day d. (Note that by converting from an absolute measure to a relative measure, we adjust for the fact that larger QSRs are more likely on average to experience larger absolute forecast errors than smaller QSRs.) The greater the value of $AmpErr_{id}$ the more inaccurate the total demand forecast on that day (i.e., the more actual demand deviates from expected demand).

4.3.3. Relative Error in Demand Pattern To measure the relative error in forecasting demand pattern, $PattErr_{id}$, we compute the Spearman rank-order correlation (Myers et al. 2013) between forecasted and observed demand. In effect, this serves to rank order the forecasted and actual sales within a day, and then determines the degree of correlation between the rank-orders.

By defining relative error in forecasting demand pattern as a function of the rank, we are abstracting away from the *amplitude* aspect of the forecast error. This enables us to separate the amplitude effect from the pattern effect, and it captures the extent to which the demand model is able to accurately identify peaks and troughs in demand within a day. This distinction can be seen in Table 2, which shows that the correlation between these two error measures is small, taking value -0.082

4.4. 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 total sales made through the platform and total 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 aforementioned idiosyncratic and unobservable

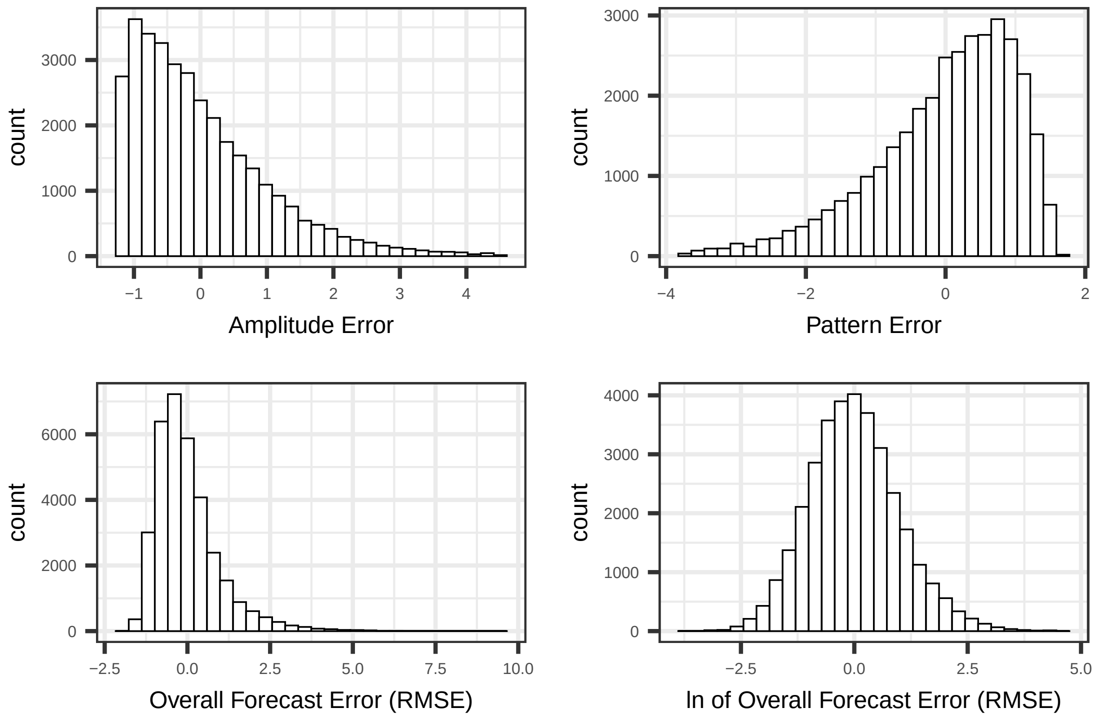

Figure 1 **Histograms of standardized measures of forecast error**

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 would therefore 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, i.e.,

$$Dep_{id} = \frac{\sum_{d \in T_d} PlatformShare_{id}}{|T_d|}. \tag{5}$$

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. (5) 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.[9]

In essence, Dep_{id} is a moving average of daily share of demand from platforms, and it 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., 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 $w = 14$.

4.5. Summary Statistics

Summary statistics and correlations between the key variables outlined in Sections 4.3–4.4 are given in Table 2. In order to maintain anonymity for the QSR chain, all summary statistics are reported on standardized variables – i.e., after subtracting the mean and dividing by the standard deviation. However, statistical analysis is performed and results reported for the 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 of interest, the descriptive statistics show that all variables appear to be well-behaved except for $OverallErr_{id}$, 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 $OverallErr_{id}$ prior to modeling.

[9] Note 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 A.

Table 2 Descriptive statistics and correlations for key variables.

Panel A: Descriptive Statistics – Standardized Variables

	Mean	Median	Max	Min	St.Dev Overall	St.Dev Between	St.Dev Within
$Sales_{id}$	0	-0.22	15.37	-1.66	1	0.86	0.60
$Diners_{id}$	0	-0.22	12.06	-1.64	1	0.83	0.64
$OverallErr_{id}$	0	-0.21	9.65	-1.85	1	0.85	0.88
$AmpErr_{id}$	0	-0.24	4.52	-1.23	1	0.86	0.97
$PattErr_{id}$	0	0.20	1.71	-3.71	1	0.98	0.89
$PlatformShare_{id}$	0	-0.17	5.43	-1.26	1	0.79	0.66
Dep_{id}	0	-0.08	3.67	-1.48	1	0.87	0.50

Panel B: Correlations

	$Sales_{id}$	$Diners_{id}$	$OverallErr_{id}$	$AmpErr_{id}$	$PattErr_{id}$	$PlatformShare_{id}$
$Diners_{id}$	0.957∗∗∗	–				
$OverallErr_{id}$	0.603∗∗∗	0.619∗∗∗	–			
$AmpErr_{id}$	-0.123∗∗∗	-0.099∗∗∗	0.293∗∗∗	–		
$PattErr_{id}$	0.283∗∗∗	0.273∗∗∗	-0.059∗∗∗	-0.082∗∗∗	–	
$PlatformShare_{id}$	-0.048∗∗∗	0.03∗∗∗	-0.033∗∗∗	0.006	0.002	–
Dep_{id}	-0.029∗∗∗	0.052∗∗∗	0.004	-0.011∗	0.01∗	0.821∗∗∗

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.

Also note 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 (cf. Table 1). Therefore, to isolate the effects of interest, we next outline the statistical modeling approach taken in this paper.

5. Models and Results

In this section we describe the estimation approach used to answer Empirical Questions 1–4, then present results and robustness.

5.1. Econometric Specification

Our empirical strategy takes advantage of 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 amplitude, demand pattern, and overall demand at each restaurant. We then perform mediation analysis to investigate the relative importance of errors in amplitude and pattern in mediating the relationship between dependence and overall forecast error.

More formally, we answer Empirical Questions 1–3 by estimating the following set of independent regression equations:

$$AmpErr_{id} = \alpha_i + \alpha_1 Dep_{id} + \boldsymbol{\alpha}_2^\top \mathbf{X}_{id} + \epsilon_{id}^{\alpha} \tag{6}$$

$$PattErr_{id} = \beta_i + \beta_1 Dep_{id} + \boldsymbol{\beta}_2^\top \mathbf{X}_{id} + \epsilon_{id}^{\beta} \tag{7}$$

$$ln(OverallErr_{id}) = \gamma_i + \gamma_1 Dep_{id} + \boldsymbol{\gamma}_2^\top \mathbf{X}_{id} + \epsilon_{id}^{\gamma} \tag{8}$$

In the above set of equations, the primary effects of interest are given by α_1, β_1 and γ_1, which correspond to the effect of an increase in platform dependence on a QSR's error in forecasting demand amplitude, demand pattern and overall demand, respectively. Meanwhile, α_i, β_i and γ_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 forecast errors. The error terms ϵ_{id}^{α}, ϵ_{id}^{β}, and ϵ_{id}^{γ} are normally distributed with a conditional mean value of zero.

The vector containing the control variables in Equations (6)–(8) 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, note that as daily demand increases, so too might the errors in our forecasts (i.e., δ_{idh} in Equation (2) may increase). As a result, the overall 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, i.e., $Scale_{id} = ln(\sum_h Demand_{idh}^F)$.

Turning to Empirical Question 4, we answer this by following the mediation approach outlined in the classic works of Judd and Kenny (1981*a*,*b*), Baron and Kenny (1986) and Kenny et al. (1998) (also see MacKinnon et al. (2007) for a lucid discussion). This requires us to estimate the following regression equations:

$$ln(OverallErr_{id}) = \theta_i + \theta_1 Dep_{id} + \theta_2 AmpErr_{id} + \boldsymbol{\theta}_3^\top \mathbf{X}_{id} + \epsilon_{id}^{\theta} \tag{9}$$

$$ln(OverallErr_{id}) = \eta_i + \eta_1 Dep_{id} + \eta_2 PattErr_{id} + \boldsymbol{\eta}_3^\top \mathbf{X}_{id} + \epsilon_{id}^{\eta}. \tag{10}$$

Equation (9) estimates the mediating effect of amplitude error on the relationship between dependence on platforms and overall forecast error. Here, θ_1 is an estimate of the average direct effect (ADE) of dependence on overall forecast error, after controlling for the mediating effect of amplitude error. The average causal mediation effect (ACME) – which is an estimate of the effect that is mediated by amplitude error – is equal to $\theta_2 \times \alpha_1$, where α_1 is estimated in Equation (6). Meanwhile, the proportion of the total effect that can be attributed to mediation by amplitude error is given by $(\theta_2 \times \alpha_1)/\gamma_1$, where γ_1 is estimated in Equation (8).

Similarly, Equation (10) estimates the mediating effect of pattern error on the relationship between dependence on platforms and overall forecast error. In this case, η_1 is an estimate of the ADE and $\eta_2 \times \beta_1$ is an estimate of ACME, while $(\eta_2 \times \beta_1)/\gamma_1$ is the proportion of the total effect that can be attributed to mediation by pattern error.

5.2. Results

Table 3 reports the estimates of coefficients from the regression equations outlined in Section 5.1. Standard errors reported are heteroskedasticity and autocorrelation consistent.

The first, second and third columns in Table 3 answer Empirical Questions 1, 2 and 3, respectively (cf. Section 3). We find that an increasing dependence on platforms has only a small effect on the error in forecasting demand amplitude (denoted by α_1 in Eq. (6)). Specifically, for every 10 percentage point (p.p.) increase in dependence on platforms, error in forecasting demand amplitude increases by 0.88 p.p. (p-value< 0.001). However, we find that increasing platform dependence has a significant impact on the error in forecasting demand pattern (denoted by β_1 in Eq. (7)): for every 10 p.p. increase in dependence on platforms, pattern correlation decreases by 2.6 p.p. (p-value < 0.001). Increasing dependence on platforms also has a significant impact on the overall error in forecasting demand (denoted by γ_1 in Eq. (8)). For every 10 p.p. increase in platform dependence, overall error in forecasting demand increases by 2.83% (p-value < 0.01).

Turning to the mediating effects, in Table 3, note that once $AmpErr$ is introduced as a mediating variable in the control structure (column four), the marginal effect of dependence on overall forecasting error falls from 28.3% to 15.9%, remaining significant only at the 10% level (p-value $= 0.065$). In the case of $PattErr$ (column five) this change is even more pronounced and the direct effect no longer remains statistically significant – falling from 28.3% to 10.8% (p-value $= 0.262$). And finally, when both the mediating variables are introduced simultaneously (column six), the marginal effect of dependence on overall forecasting error is negligibly small and statistically insignificant. This indicates that the increase in platform dependence on forecast errors can, essentially, be entirely explained by an increase in errors associated with forecasting the daily amplitude and within-day pattern of demand.

To explore this further, we follow the mediation approach described in Section 5.1. To understand mediation by $AmpErr$ we combine the coefficient estimates from columns (1), (3) and (4) in Table 3. Using this, we estimate the ACME ($\theta_2 \times \alpha_1$) to be 0.124 and the proportion of effect mediated by $AmpErr$ to be 43.9% ($(\theta_2 \times \alpha_1)/\gamma_1$). Using columns (2), (3) and (5) from Table 3 to explore mediation by $PattErr$ gives an ACME ($\eta_2 \times \beta_1$) of 0.175, while the proportion of effect mediated by $PattErr$ is estimated to be 61.8% ($\eta_2 \times \beta_1/\gamma_1$). However, strictly speaking, additional analysis is required to rigorously establish the statistical significance of these mediation effects (Singh and Fleming 2010). To that end, we use a bootstrapping approach developed by Preacher and Hayes (2004) to corroborate these mediated effect sizes and confirm their statistical significance (see Appendix B for results from the bootstrapping analysis).

Overall, these results demonstrate that increasing dependence on platform demand results in poorer demand forecasting performance across the board, with this especially negatively affecting a restaurants ability to forecast at what time demand will materialize within a day. The unpredictable nature of the timing of these arrivals is further worsened by the fact that total daily demand is also harder to forecast as restaurants increase their dependence on these platforms.

5.3. Alternative Explanations

It is seldom straightforward to establish causality through 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.

First, by taking advantage of only the within-restaurant variation in dependence and forecast error to estimate our models, we control for a wide variety of unobserved site-specific effects, such as restaurant size, neighborhood demographics, intensity of competitor activity, etc.

Second, our findings are based on the underlying degree of *dependence* each restaurant has on demand that comes through food-delivery platforms (cf. Section 4.4 and Appendix A for definitions of Dep_{id} and $PlatformShare_{id}$). In order 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. To further account for confounding by such multi-day factors, we

Table 3 Main regression results

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.088***	−0.259***	0.283***	0.159*	0.108	−0.005
	(0.023)	(0.051)	(0.096)	(0.086)	(0.096)	(0.088)
Amplitude error				1.412***		1.377***
				(0.035)		(0.034)
Pattern correlation					−0.675***	−0.645***
					(0.026)	(0.023)
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	32,695	32,695	32,695	32,695
R^2	0.076	0.049	0.350	0.506	0.440	0.588
Adjusted R^2	0.072	0.045	0.348	0.504	0.438	0.586

Note: *p<0.1; **p<0.05; ***p<0.01
Robust standard errors in parantheses below each coefficient.

also (i) perform robustness checks by excluding additional days around the focal day when calculating dependence and (ii) estimate a model that includes the lag of the dependent variable on the right hand side (see Section 5.4 for more details).

Third, we use a comprehensive range of control variables (cf. Table 1) to account for confounding factors. These factors, such as promotional campaigns and extended holiday periods, can have an impact on both *dependence* and *forecast error*. By explicitly including them in the control structure of our linear regressions, we partial out their impact on the relationship between *dependence* and *forecast error*.

Finally, it is possible that the increase in forecast error is due to an increase in the total demand – which could be due to either an increasing dependence on platforms or some other unknown variable that is increasing over time. We control for this in two ways: we include trend as a control for any omitted variables that are increasing with time, and we include aggregated day-level forecasts as a control to account for the scale of the demand for which forecasts are made.

5.4. Robustness Tests

We 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, we use alternative specifications for the demand model from which we extracted the fits for computing forecast error (see the end of Section 4.2.2 for more details); we present these results in Appendix C. More specifically, in Appendix C.2, we use a simpler model that includes only staffing data and restaurant-specific factors as features, ignoring other features listed in Table 1. Additionally, in Appendix C.3, we re-run our analysis with a more sophisticated model that uses an ensemble of machine learning methods.

We also perform additional checks on the results discussed in Table 3, where we use RMSE to measure overall forecast error. We repeat our analysis using Mean Absolute Error (MAE) to measure overall forecast error; see Appendix D for results of this re-estimation, which are similar to our original findings. We also repeat our analysis employing a range of window sizes to measure dependence on platforms and report results for 10- and 20-day windows in Appendix E.1 and Appendix E.2, respectively.

When computing the moving average of $PlatformShare$ to measure Dep_{id}, we eliminate confounding factors by disregarding the focal day d (cf. Appendix A for full details of the moving average specification). However, eliminating the focal day itself does not account for any potential confounding factors that can have an influence beyond the focal day. To ascertain the robustness of our results, we compute variants of Dep_{id} that also disregard $PlatformShare$ from a few days surrounding (and including) the focal day d. We discuss the details of this modified moving average specification for Dep_{id} in Appendix F. In Appendix F.1, we discuss results with Dep_{id} measured as a moving average of $PlatformShare$, but disregarding the focal day and ± 1 day surrounding the focal day. In Appendices F.2 and F.3, we disregard the focal day and, respectively, ± 2 and ± 3 days surrounding it.

Another technique that can eliminate, or at least mitigate, the impact of confounding factors is to use a proxy variable that can account for the unobserved variable. Woolridge et al. (2002) note that "often the outcome of the dependent variable from an earlier time period can be a useful proxy variable" (p. 66) to deal with the bias that arises from an omitted variable. This technique has been used in previous empirical studies – for instance, see Heckman and Borjas (1980), Nerkar and Paruchuri (2005), Gokpinar et al. (2010). In Appendix G, we present results from regressions that include the first lag of forecast error as a covariate (in addition to those features presented in Table 1).

Overall, regardless of the demand model specified or the way we measure dependence and forecast error, we consistently find that as the QSR increases its dependence on third-party platforms, its overall error in forecasting demand increases. We also find that increasing dependence has only a small and positive impact on error in forecasting demand amplitude, or sometimes an insignificant impact (at 10% p-value level). Whatever the case, accuracy in forecasting demand patterns consistently deteriorates as QSRs become more dependent on platforms.[10] And finally, models that account for changes in both demand amplitude and demand patterns consistently have a statistically insignificant residual relationship between dependence and overall error in demand forecasts.

6. Conclusions and Discussion

The food delivery market is expected to reach a total value significantly over US$100 billion in the next two to three years, and the COVID-19 pandemic is, if anything, only accelerating its growth. Powered by app-based platforms that connect restaurants, drivers/riders and consumers in real time, it is changing the food and beverage industry landscape. The high popularity of these delivery platforms among consumers has driven restaurants to increasingly depend on them to drive their sales growth. However, restaurant managers need to be aware of the potential adverse effects of this dependence.

In this paper, we conduct a rigorous analysis to establish, first, that as QSRs' dependence on third-party food-delivery platforms increases, their ability to accurately forecast their demand decreases. This ability is critical for efficient planning in restaurant operations. It is especially important in the context of QSRs, where speed of service is a key competitive advantage and cooked food has a short shelf-life, leading to high potential for food waste. Second, we show that the increase in forecast error can be primarily attributed to errors in forecasting demand *pattern*, as opposed to forecasting demand *amplitude*.

Our work has important implications for the nascent literature examining the relationship between new, app-based marketplaces and traditional businesses (e.g., Feldman et al. 2018, Chen, Hu and Wang 2019). We empirically establish that demand coming from food-delivery platforms changes the characteristics of the overall demand at a restaurant. Going forward, this dynamic needs to be properly accounted for in models that study the relationship between delivery platforms and restaurants.

[10] This result is confirmed by a survey-based study of anecdotal evidence from restaurateurs in the UK. (Smithers 2020)

From a managerial perspective, our work adds several rigorously-tested insights to the conversation among restaurateurs on the unintended consequences of third-party food-delivery platforms on their businesses. Since an increase in forecast error leads to an increase in a restaurant's operational costs, managers must account for this when considering a relationship with third-party delivery platforms and offset that against any anticipated increase in their revenues. The overall impact will be context specific and our paper provides a methodology that restaurant managers can follow to measure these trade-offs and to, at regular intervals, re-assess their engagement with their platform partners.

To improve adoption, the platform itself may be able to help restaurants by offering value-added services that help their client restaurants to anticipate demand patterns and lower the costs of adoption. Since the platform has access to real-time data on platform demand for all of the restaurants that are part of their network, such forecasts will likely be more accurate than those that the restaurant can produce using only their own data. For example, restaurants and platforms might negotiate information sharing agreements that can help restaurants improve their forecasting and, in general, operational efficiency. Advantages of information sharing in supply chains are well understood (for instance, cf. Lee and Whang 2000), and undertaken by companies such as Amazon (Levy 2019) and Airbnb (Novet 2015).

Extrapolating from our results, we also hypothesize that demand at so-called "virtual kitchens"[11] will be much harder to forecast than demand at a traditional restaurant. This implies that margins on cloud kitchens will be lower than expected due to the additional costs of decreased forecast accuracy. This implication is highly relevant for managers and shareholders of such businesses. In fact, recent news reports indicate that firms that have invested in cloud kitchens are now realizing the "high volatility" associated with those businesses and are scaling down their investments (Singh 2020).

Finally, our results suggest that changes in forecast accuracy in demand amplitude and demand patterns explain almost all the deterioration in overall forecast accuracy. This suggests that restaurants, in responding to this, can focus their attention on just two strategies: one, improving forecast accuracy for total daily demand; two, improving forecast accuracy for distribution of demand within the day. Further analysis suggest that restaurants should

[11] Also known as cloud, ghost or dark kitchens, these kitchens only serve delivery customers and do not have any seating or direct customer-facing component.

rather focus their attention on the latter. To do this, they could take advantage of modern IS-enabled capabilities to offer personalized, geolocation-based and time-of-day-based recommendations and incentives to customers. There are some indications that customers may be amenable to this; for instance, Susskind et al. (2004) show that consumers are willing to change their dining times when proper incentives are offered. In this day and age, restaurants cannot avoid depending on third-party food-delivery platforms, but our results suggest that they should endeavor to regulate their demand patterns in order to take full advantage of this growing phenomenon.

References

Agatz, N. A., Fleischmann, M. and Van Nunen, J. A. (2008), 'E-fulfillment and multi-channel distribution–a review', *European journal of operational research* **187**(2), 339–356.

Ahuja, V., Alvarez, C. and Staats, B. R. (2019), 'How continuity in service impacts process variability: Evidence from a primary care setting', *Available at SSRN 3440355* .

Akturk, M. S., Ketzenberg, M. and Heim, G. R. (2018), 'Assessing impacts of introducing ship-to-store service on sales and returns in omnichannel retailing: A data analytics study', *Journal of Operations Management* **61**, 15–45.

Allon, G., Federgruen, A. and 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.

Andaleeb, S. S. and Conway, C. (2006), 'Customer satisfaction in the restaurant industry: an examination of the transaction-specific model', *Journal of services marketing* .

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

Awasthi, P. (2020), 'Diners have inhibition to visit their favourite restaurants post lockdown: Survey', *The Hindu Business Line* .
URL: *https://www.thehindubusinessline.com/news/variety/diners-have-inhibition-to-visit-their-favourite-restaurants-post-lockdown-survey/article31640412.ece*

Bakos, J. Y. (1991), 'A strategic analysis of electronic marketplaces', *MIS quarterly* pp. 295–310.

Bakos, J. Y. (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. and Riquelme, C. (2016), 'Dynamic pricing in ridesharing platforms', *ACM SIGecom Exchanges* **15**(1), 65–70.

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

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

Baron, R. M. and Kenny, D. A. (1986), 'The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations.', *Journal of personality and social psychology* **51**(6), 1173.

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

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

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

Cachon, G. P., Daniels, K. M. and 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, G. P. and Terwiesch, C. (2016), *Operations Management*, McGraw-Hill.

Cagle, S. (2020), ''fees are murder': delivery apps profit as restaurants forced to close doors', *The Guardian* .
URL: *https://www.theguardian.com/world/2020/apr/03/delivery-app-restaurants-coronavirus-california*

Chandrashekar, T. S., Narahari, Y., Rosa, C. H., Kulkarni, D. M., Tew, J. D. and Dayama, P. (2007), 'Auction-based mechanisms for electronic procurement', *IEEE Transactions on Automation Science and Engineering* **4**(3), 297–321.

Chen, M., Hu, M. and Wang, J. (2019), 'Food delivery service and restaurant: Friend or foe?', *Available at SSRN 3469971* .

Chen, R. R., Roundy, R. O., Zhang, R. Q. and Janakiraman, G. (2005), 'Efficient auction mechanisms for supply chain procurement', *Management Science* **51**(3), 467–482.

Chen, T. and Guestrin, C. (2016), Xgboost: A scalable tree boosting system, *in* 'Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining', pp. 785–794.

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

Chiappetta, M. (2020), 'Uber eats demand soars due to covid-19 crisis', *Forbes* .
URL: *https://www.forbes.com/sites/marcochiappetta/2020/03/25/uber-eats-demand-soars-due-to-covid-19-crisis*

De Vries, J., Roy, D. and De Koster, R. (2018), 'Worth the wait? how restaurant waiting time influences customer behavior and revenue', *Journal of Operations Management* .

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, A. E. and Swinney, R. (2018), 'Service delivery platforms: Pricing and revenue implications', *Available at SSRN 3258739* .

First, D. (2020), 'When will diners be ready to return to restaurants? here's what 1,000 globe readers said', *Boston Globe* .
URL: *https://www.bostonglobe.com/2020/05/26/lifestyle/when-will-diners-be-ready-return-restaurants*

Freeman, M., Robinson, S. and Scholtes, S. (2017), 'Gatekeeping under congestion: An empirical study of referral errors in the emergency department'.

Friedman, J. H. (2001), 'Greedy function approximation: a gradient boosting machine', *Annals of statistics* pp. 1189–1232.

Gallino, S., Moreno, A. and Stamatopoulos, I. (2017), 'Channel integration, sales dispersion, and inventory management', *Management Science* **63**(9), 2813–2831.

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

Gokpinar, B., Hopp, W. J. and Iravani, S. M. (2010), 'The impact of misalignment of organizational structure and product architecture on quality in complex product development', *Management science* **56**(3), 468–484.

Gunders, D. (2017), *Wasted: How America is losing up to 40 percent of its food from farm to fork to landfill.*

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

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

Heckman, J. J. and Borjas, G. J. (1980), 'Does unemployment cause future unemployment? definitions, questions and answers from a continuous time model of heterogeneity and state dependence', *Economica* **47**(187), 247–283.

Isaac, M. and 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*

Judd, C. M. and Kenny, D. A. (1981*a*), *Estimating the effects of social intervention*, CUP Archive.

Judd, C. M. and Kenny, D. A. (1981*b*), 'Process analysis: Estimating mediation in treatment evaluations', *Evaluation review* **5**(5), 602–619.

Kapoor, R. and 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.

Kenny, D. A., Kashy, D. A., Bolger, N. et al. (1998), 'Data analysis in social psychology', *The handbook of social psychology* **1**(4), 233–265.

Khan, M. (2019), 'The rush to offer food deliveries is leaving franchisees at the crossroads', *Forbes* .
URL: *https://www.forbes.com/sites/mahmoodkhan1/2019/08/10/the-rush-to-offer-food-deliveries-is-leaving-franchisees-at-the-crossroads*

Klein, D. (2019), 'Fighting for share in the usd 16.6 billion delivery app market', *QSR Magazine* .
URL: *https://www.qsrmagazine.com/technology/fighting-share-166-billion-delivery-app-market*

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

Lee, H. L. and Whang, S. (2000), 'Information sharing in a supply chain', *International Journal of Manufacturing Technology and Management* **1**(1), 79–93.

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

Levy, N. (2019), 'Amazon unveils small business academy to help third-party sellers bolster e-commerce presence', *Geekwire* .
URL: *https://www.geekwire.com/2019/amazon-unveils-small-business-academy-help-third-party-sellers-bolster-e-commerce-presence/*

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

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

Lyu, G., Cheung, W. C., Teo, C.-P. and Wang, H. (2019), 'Multi-objective online ride-matching', *Available at SSRN 3356823* .

MacKinnon, D. P., Fairchild, A. J. and Fritz, M. S. (2007), 'Mediation analysis', *Annu. Rev. Psychol.* **58**, 593–614.

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, T. W., Yates, J. and Benjamin, R. I. (1987), 'Electronic markets and electronic hierarchies', *Communications of the ACM* **30**(6), 484–497.

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

Myers, J. L., Well, A. D. and Lorch Jr, R. F. (2013), *Research design and statistical analysis*, Routledge.

Nerkar, A. and Paruchuri, S. (2005), 'Evolution of r&d capabilities: The role of knowledge networks within a firm', *Management science* **51**(5), 771–785.

Novet, J. (2015), 'Airbnb is introducing pricing tips for hosts', *VentureBeat* .
URL: *https://venturebeat.com/2015/06/04/airbnb-is-introducing-pricing-tips-for-hosts/*

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

Pereira, D. X. R. (2018), 'Going zero waste in canteens: Exploring food demand using data analytics'.

Pittman, R. (2019), 'Fast-food drive thrus got 20 seconds slower in 2019', *QSR Magazine* .
URL: *https://www.qsrmagazine.com/drive-thru/fast-food-drive-thrus-got-20-seconds-slower-2019*

Preacher, K. J. and Hayes, A. F. (2004), 'Spss and sas procedures for estimating indirect effects in simple mediation models', *Behavior research methods, instruments, & computers* **36**(4), 717–731.

ReFED, A. (2016), 'A roadmap to reduce us food waste by 20 percent'.

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

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

Singh, J. and Fleming, L. (2010), 'Lone inventors as sources of breakthroughs: Myth or reality?', *Management science* **56**(1), 41–56.

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/*

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

Smithers, R. (2020), 'Uk takeaway food waste rises during coronavirus lockdown', *The Guardian* .
URL: *https://www.theguardian.com/environment/2020/may/13/uk-takeaway-food-waste-rises-during-coronavirus-lockdown*

Soh, C., Markus, M. L. and Goh, K. H. (2006), 'Electronic marketplaces and price transparency: strategy, information technology, and success', *MIS quarterly* pp. 705–723.

Standing, C., Love, P. E., Stockdale, R. and 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. and Love, P. E. (2010), 'A review of research on e-marketplaces 1997–2008', *Decision Support Systems* **49**(1), 41–51.

Susskind, A. M., Reynolds, D. and Tsuchiya, E. (2004), 'An evaluation of guests' preferred incentives to shift time-variable demand in restaurants', *Cornell Hotel and Restaurant Administration Quarterly* **45**(1), 68–84.

Swaminathan, J. M. and Lee, H. L. (2003), 'Design for postponement', *Handbooks in Operations Research and Management Science* **11**, 199–226.

Tan, T. F. and Netessine, S. (2014), 'When does the devil make work? an empirical study of the impact of workload on worker productivity', *Management Science* **60**(6), 1574–1593.

Thompson, G. M. (1998), 'Labor scheduling, part 2: Knowing how many on-duty employees to schedule', *Cornell Hotel and Restaurant Administration Quarterly* **39**(6), 26–37.

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

Tingley, D., Yamamoto, T., Hirose, K., Keele, L. and Imai, K. (2014), 'Mediation: R package for causal mediation analysis'.

Tkacik, M. (2020), 'Restaurants are barely surviving. delivery apps will kill them.', *The Washington Post* . **URL:** *https://www.washingtonpost.com/outlook/2020/05/29/delivery-apps-restaurants-coronavirus/*

Ülkü, S., Hydock, C. and Cui, S. (2019), 'Making the wait worthwhile: Experiments on the effect of queueing on consumption', *Management Science* .

Wang, S., Zheng, S., Xu, L., Li, D. and 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/*

Woolridge, J. M. et al. (2002), 'Econometric analysis of cross section and panel data'.

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

Zhang, J., Farris, P. W., Irvin, J. W., Kushwaha, T., Steenburgh, T. J. and Weitz, B. A. (2010), 'Crafting integrated multichannel retailing strategies', *Journal of Interactive Marketing* **24**(2), 168–180.

Appendix A: 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 4.4, $PlatformShare_{id}$ 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).

$$PlatformShare_{id} = Dep_{id} + noise_{id} \tag{11}$$

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 $PlatformShare_{id}$.

Since the impact of these promotional campaigns may be non-random, 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:

$$PlatformShare_{it} = \lambda_{id} + \gamma_{id} \times banner.binary_{it} + \Delta_{it} \tag{12}$$

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). In Appendix E we present results with alternate window sizes.

Appendix B: Mediation Analysis with Bootstrapped Samples

We test the statistical significance of the mediation effects of forecasting error in demand pattern and demand amplitude via non-parametric bootstrapping analysis as outlined in Preacher and Hayes (2004). We generate confidence intervals for this analysis from 1500 simulations, using the `mediation` package in `R` statistical software developed by Tingley et al. (2014). Table 4 and Table 5 report non-parametric bootstrap confidence intervals for the mediation effects we present in Section 5.2.

Table 4 Nonparametric Bootstrap Confidence Intervals for Mediation by Demand Amplitude

	Estimate	95% CI Lower	95% CI Upper	p-value	
ACME	0.122	0.058	0.187	0	***
ADE	0.158	0.039	0.269	0.012	*
Total Effect	0.280	0.142	0.407	0	***
Prop. Mediated	0.436	0.247	0.753	0	***

Table 5 Nonparametric Bootstrap Confidence Intervals for Mediation by Demand Pattern

	Estimate	95% CI Lower	95% CI Upper	p-value	
ACME	0.175	0.122	0.221	0	***
ADE	0.105	-0.016	0.228	0.088	.
Total Effect	0.280	0.148	0.407	0	***
Prop. Mediated	0.624	0.401	1.104	0	***

Appendix C: Alternative Demand Models

C.1. Number of diners served as the measure of demand

In Section 4.2 we mention that restaurant demand can also be measured by number of diners served. In this section, we repeat the analysis from our main manuscript using number of diners served as the measure of demand, and our results and insights still hold.

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.

Table 6 Results with no. of diners served as measure of demand

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.045	-0.283^{***}	0.228^{**}	0.163^{**}	0.037	-0.018
	(0.030)	(0.052)	(0.095)	(0.078)	(0.101)	(0.086)
Amplitude error				1.474^{***}		1.439^{***}
				(0.038)		(0.035)
Pattern correlation					-0.676^{***}	-0.642^{***}
					(0.027)	(0.024)
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	32,699	32,699	32,699	32,699	32,699	32,699
R^2	0.086	0.050	0.390	0.547	0.466	0.616
Adjusted R^2	0.083	0.047	0.387	0.545	0.464	0.614

Note: *p<0.1; **p<0.05; ***p<0.01
Robust standard errors in parantheses below each coefficient.

C.2. Man-hours only model

In this section, we present results when we recreate our analysis using forecasts from a simpler model including only staffing data and restaurant-specific factors as features and ignoring the other variables listed in Table 1. This model provides an out-of-sample R-square of 42.6%, and it captures a scenario where QSR managers are less sophisticated and rely primarily on intuition and experience to estimate demand rather than using statistical models.

Table 7 Results with forecasts from a simpler model

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.101	-0.267^{*}	0.251^{*}	0.211	0.177	0.134
	(0.094)	(0.138)	(0.132)	(0.128)	(0.124)	(0.123)
Amplitude error				0.401^{***}		0.403^{***}
				(0.098)		(0.097)
Pattern correlation					-0.279^{***}	-0.286^{***}
					(0.016)	(0.013)
Day of the week	Yes	Yes	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes	Yes	Yes
Week of the month	Yes	Yes	Yes	Yes	Yes	Yes
Public holiday	Yes	Yes	Yes	Yes	Yes	Yes
Religious 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	33,302	33,302	33,302	33,302	33,302	33,302
R^2	0.120	0.123	0.182	0.351	0.231	0.402
Adjusted R^2	0.116	0.120	0.179	0.349	0.228	0.400

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Robust standard errors in parantheses below each coefficient.

C.3. Forecasts from a stacked demand model

In this section, we present results when we recreate our analysis using forecasts from a more sophisticated model using machine learning and ensemble methods to combine the forecasts of several algorithms (e.g., gradient boosting and random forests) into a single aggregate forecast.

We use the `h2o.automl` function from the `h2o` package in `R` for model training and selection. `h2o.automl` can be used for automatically training and tuning machine learning models with minimal input from the user (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 several models. The stacked ensemble model generated forecasts by aggregating forecasts from constituent individual 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%.

Table 8 Results with forecasts from a more sophisticated model

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.016	-0.265^{***}	0.326^{***}	0.294^{***}	0.110	0.100
	(0.019)	(0.051)	(0.089)	(0.069)	(0.082)	(0.070)
Amplitude error				2.022^{***}		1.894^{***}
				(0.064)		(0.057)
Pattern correlation					-0.815^{***}	-0.740^{***}
					(0.035)	(0.030)
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	33,308	33,308	33,308	33,308	33,308	33,308
R^2	0.028	0.080	0.240	0.398	0.352	0.489
Adjusted R^2	0.025	0.077	0.237	0.395	0.350	0.487

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Robust standard errors in parantheses below each coefficient.

Appendix D: Regressions with MAE

As we discuss in Section 4.3.1, Mean Absolute Error (MAE) is another metric commonly used to measure the accuracy of forecasts. We find that our results are robust to specifications with this measure.

Table 9 Results with MAE as measure of forecast error

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.088^{***}	-0.259^{***}	0.273^{***}	0.148^{*}	0.106	−0.009
	(0.023)	(0.051)	(0.099)	(0.089)	(0.102)	(0.092)
Amplitude error				1.430^{***}		1.396^{***}
				(0.038)		(0.036)
Pattern correlation					-0.648^{***}	-0.618^{***}
					(0.024)	(0.021)
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	32,695	32,695	32,695	32,695
R^2	0.076	0.049	0.369	0.537	0.456	0.615
Adjusted R^2	0.072	0.045	0.367	0.535	0.454	0.614

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Robust standard errors in parantheses below each coefficient.

Appendix E: Alternative Window Sizes for Computing Dep_{id}

E.1. 10 days

In our main results, we chose a 28-day window to run the regressions. In this section, we show that our results are robust to other window size specifications. For instance, we re-run the regressions using a ±10-day window, which is approximately equivalent to averaging over 1.5 weeks in both directions: past and future.

Table 10 Results with Dependence defined over a three weeks moving window

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.038^{**}	-0.201^{***}	0.228^{***}	0.174^{**}	0.092	0.045
	(0.018)	(0.047)	(0.084)	(0.076)	(0.088)	(0.080)
Amplitude error				1.414^{***}		1.380^{***}
				(0.036)		(0.034)
Pattern correlation					-0.675^{***}	-0.645^{***}
					(0.027)	(0.023)
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	32,874	32,874	32,874	32,874	32,874	32,874
R^2	0.075	0.048	0.349	0.506	0.439	0.588
Adjusted R^2	0.071	0.044	0.347	0.504	0.437	0.586

Note: *p<0.1; **p<0.05; ***p<0.01

Robust standard errors in parantheses below each coefficient.

E.2. 20 days

Similar to the previous table, table Table 11 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 11 Results with Dependence defined over a six weeks moving window

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.068^{**}	-0.273^{***}	0.288^{**}	0.192^{*}	0.104	0.018
	(0.030)	(0.056)	(0.119)	(0.099)	(0.115)	(0.099)
Amplitude error				1.411^{***}		1.375^{***}
				(0.036)		(0.034)
Pattern correlation					-0.676^{***}	-0.645^{***}
					(0.026)	(0.023)
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	32,296	32,296	32,296	32,296	32,296	32,296
R^2	0.076	0.049	0.352	0.507	0.441	0.589
Adjusted R^2	0.072	0.045	0.349	0.506	0.439	0.587

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Robust standard errors in parantheses below each coefficient.

Appendix F: Alternative Definitions of Dependence

In Appendix A 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).

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. 12 over $t \in [d-w, d+w] \wedge t \notin [d-l, d+l]$, where $w = 14$ and $l \in \{1, 2, 3\}$.

F.1. Moving average with $w = 14$ and $l = 1$

Table 12 Results with an alternative definition of Dependence

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.095***	−0.251***	0.298***	0.165*	0.129	0.006
	(0.025)	(0.050)	(0.098)	(0.087)	(0.097)	(0.088)
Amplitude error				1.410***		1.375***
				(0.036)		(0.034)
Pattern correlation					−0.674***	−0.645***
					(0.026)	(0.023)
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	32,486	32,486	32,486	32,486	32,486	32,486
R^2	0.076	0.049	0.351	0.506	0.440	0.588
Adjusted R^2	0.072	0.045	0.349	0.504	0.438	0.586

Note: *p<0.1; **p<0.05; ***p<0.01
Robust standard errors in parantheses below each coefficient.

F.2. Moving average with $w = 14$ and $l = 2$

Table 13 Results with an alternative definition of Dependence

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.106^{***}	-0.247^{***}	0.321^{***}	0.172^{**}	0.155	0.017
	(0.027)	(0.049)	(0.096)	(0.084)	(0.095)	(0.085)
Amplitude error				1.410^{***}		1.375^{***}
				(0.036)		(0.034)
Pattern correlation					-0.674^{***}	-0.645^{***}
					(0.026)	(0.023)
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	32,488	32,488	32,488	32,488	32,488	32,488
R^2	0.076	0.049	0.351	0.506	0.440	0.588
Adjusted R^2	0.073	0.045	0.349	0.504	0.438	0.586

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Robust standard errors in parantheses below each coefficient.

F.3. Moving Average with $w = 14$ and $l = 3$

Table 14 Results with an alternative definition of Dependence

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.104^{***}	-0.253^{***}	0.338^{***}	0.191^{**}	0.167^{*}	0.031
	(0.028)	(0.047)	(0.097)	(0.085)	(0.094)	(0.084)
Amplitude error				1.411^{***}		1.375^{***}
				(0.036)		(0.034)
Pattern correlation					-0.675^{***}	-0.645^{***}
					(0.026)	(0.023)
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	32,278	32,278	32,278	32,278	32,278	32,278
R^2	0.076	0.049	0.352	0.507	0.441	0.589
Adjusted R^2	0.073	0.045	0.349	0.506	0.439	0.587

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Robust standard errors in parantheses below each coefficient.

Appendix G: 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 15 Results from models that include first lag of DV a control

	Amplitude error	Pattern correlation	Overall forecast error			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependence	0.074***	−0.248***	0.224**	0.118	0.052	−0.045
	(0.020)	(0.048)	(0.088)	(0.084)	(0.091)	(0.087)
Amplitude error				1.399***		1.364***
				(0.034)		(0.033)
Pattern correlation					−0.679***	−0.650***
					(0.024)	(0.022)
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
First lag Amplitude Error	Yes	No	No	No	No	No
First lag Pattern Correlation	No	Yes	No	No	No	No
First lag Overall Forecast Error	No	No	Yes	Yes	Yes	Yes
Observations	32,170	32,170	32,170	32,170	32,170	32,170
R^2	0.093	0.049	0.355	0.506	0.446	0.590
Adjusted R^2	0.090	0.045	0.353	0.504	0.444	0.588

Note: *p<0.1; **p<0.05; ***p<0.01
Robust standard errors in parantheses below each coefficient.