

The Behavioral, Operational and Sustainability Effects of Fast Delivery in E-Commerce

Chenshan Hu¹, Xiaoyang Long², Jiankun Sun³, Dennis J. Zhang¹

1. Olin Business School, Washington University in St. Louis, St. Louis, MO, USA

2. Wisconsin School of Business, University of Wisconsin-Madison, Madison, WI, USA

3. Imperial College Business School, Imperial College London, London, UK

It is well-established that faster delivery in e-commerce increases consumer demand. However, the impact of faster delivery on *how* consumers place orders (their order frequency and basket sizes) and the subsequent operational and environmental implications are not known. To investigate these questions, we leverage a quasi-experiment involving the opening of a new local warehouse by Alibaba Group, which led to a half-day improvement in the delivery speed for local orders. Through a difference-in-differences analysis, we find that the delivery speed improvement not only increased consumers' monthly purchasing amount by 11.05%, but also increased monthly order frequency by a higher percentage (i.e., 15.00%) and reduced the average order basket size by 1.14%. These results collectively suggest that with faster delivery, consumers purchase more on the platform but do so in *more frequent* and *smaller* orders, which implies more packaging and transportation costs for each unit of sales. We explore and identify two mechanisms contributing to the phenomenon: *order-splitting* of existing consumer demand and *category expansion* into new demand. Based on these results, we conduct a detailed calibration using both public and company-specific data to estimate the increase in the platform's operational costs and carbon emissions due to faster delivery. Our results suggest that any assessment of faster delivery initiatives need to take into account potential changes in consumers ordering behavior. As delivery becomes faster, to maintain platforms' profit margins, parallel initiatives such as incentives for order consolidation (either by warehouses or by consumers) may become more relevant.

Key words: e-commerce logistics, sustainable operations, platform operations, delivery speed, carbon emissions, shipping policy.

1. Introduction

E-commerce is the fastest-growing channel in retail. According to the U.S. Department of Commerce, e-commerce sales in the United States from 2022 accounted for 14.6% of total retail sales and was estimated at \$1,034.1 billion, an increase of 7.7% from 2021 (U.S. Census Bureau 2023). In order to drive such rapid sales explosions and to stay competitive, e-commerce companies around the world (e.g., Amazon.com, Alibaba group, etc.) are continuously improving their logistics services and, in particular, improving their delivery speed. For example, Amazon has invested heavily in building its own fleet to halve its standard delivery time—from two-day delivery to one-day delivery in most U.S. regions—on a wide range of products (*Guardian* 2020). The company also

launched a new \$1 billion venture fund to invest in companies working in the consumer fulfillment, logistics and supply chain sectors (Bloomberg 2022). In France, fast e-commerce parcel deliveries (less than 48 hours) doubled over five years as a share of the total, from about 15% to 30% (Oliver Wyman 2023). Chinese e-commerce giants, Alibaba and JD.com, have started to provide half-day delivery, or even one-hour delivery, in major cities in China by leveraging efficient logistics solutions and focusing on the last mile.

The e-commerce companies' race to improve their delivery speed stems from the belief that faster delivery leads to more sales—a belief well-supported by research (Fisher et al. 2019, Lim et al. 2023). However, a less examined but equally important question is how faster delivery changes other aspects of consumers' ordering behavior, such as their order frequency and basket size. The latter two elements are critical determinants of the operational costs (and hence the profitability) of these companies. For instance, if faster delivery induces consumers to order more online but to do so via more consolidated orders—maybe maintaining or even reducing their order frequency, then the delivery speed improvement could potentially enhance sales without increasing the operational costs related to delivery (such as transportation and packaging costs). As such, the companies' profit margins for each sale would increase, further enhancing the return from delivery speed investments. On the other hand, if faster delivery induces consumers to not only order more but to do so in smaller orders, then the opposite would hold and delivery-related costs would become more prominent in the company's overall cost structure, reducing the profitability of each sale.

Aside from its implications for profitability, a deeper understanding of how faster delivery affects consumer ordering behavior is also critical from both an operational logistics and an environmental sustainability perspective. From an operational standpoint, as e-commerce platforms continuously invest in faster delivery, accurate forecasting of fulfillment volumes and per-order costs becomes essential for capacity planning, warehouse management, and last-mile delivery operations. If faster delivery systematically shifts consumers toward more frequent, smaller orders, then platforms that fail to account for this behavioral change may face unanticipated strain on their fulfillment infrastructure and underestimate the resources required to serve the same level of demand. From an environmental standpoint, this question is equally pressing. Amid recent global calls for action on curbing carbon emissions, many companies have set carbon reduction targets, such as Amazon's pledge to achieve net-zero carbon emissions by 2040.¹ Since a significant portion of these emissions can be attributed to transportation and packaging,² any change in consumer ordering behavior

¹ <https://www.aboutamazon.com/planet/climate-pledge>

² A recent report estimates that packaging and transportation/logistics account for, respectively, 45% and 15% of carbon emissions in e-commerce (Fernández Briseño et al. 2020).

could entail a significant impact on companies' carbon reports and hence induce changing priorities among competing carbon-reduction initiatives.

In this paper, we empirically examine the impact of faster delivery on consumer ordering behavior and discuss its subsequent effect on the operational costs and carbon footprint associated with transportation and packaging. We conduct our empirical analyses by collaborating with the Alibaba Group, a leading Chinese e-commerce platform founded in 1999. Specifically, we collect data from Tmall supermarket, a large business-to-consumer (B2C) online retail platform operated by Alibaba Group. Throughout the last couple of years, Alibaba has improved its delivery speed mostly by opening local warehouses to store and shorten the shipping distance of its items. Utilizing the opening of a new warehouse in Zhengzhou, a major city in east-central China, in late August 2017, we employ a difference-in-differences (DID) approach. This method compares consumers in Zhengzhou with those in Jinan, a nearby city in a neighboring province, to estimate the average treatment effect of changes in delivery speed on consumers' subsequent purchasing behaviors. We analyze consumers' purchasing behavior in the six months before and after the warehouse opening (excluding the transition period), using a sample of existing consumers from the two cities with purchase activity in the six months prior to this study period.

We conduct several checks to ensure the internal validity of our identification strategy, thereby allowing for causal interpretations of our empirical results. First, we confirm that the warehouse opening in Zhengzhou was pre-determined by the company's logistics planning rather than driven by local sales trends and was not communicated to consumers in advance, and thus it could serve as an exogenous shock to consumers. Second, we conduct a manipulation check and confirm that the warehouse opening in Zhengzhou indeed reduced the delivery time for orders in Zhengzhou (treated city) from about 2.5 days to less than 2 days, equivalent to a 23.63% improvement in delivery speed. Third, we perform extensive tests to verify the parallel trends assumption required for the DID estimation and show that the treated and control cities exhibit parallel pre-intervention trends.

Through our DID analysis, we show that the half-day improvement in delivery speed led to a 11.05% increase in consumers' monthly purchasing amount, and a 15.00% increase in consumers' monthly order frequency. We note that the percentage increase in monthly order frequency is larger than that in monthly purchasing amount, which suggests that as delivery speed improves, consumers may be breaking their orders into smaller amounts. We further validate this by analyzing the impact of delivery speed improvement on consumers' basket size. We find that compared to the control city, the average basket size in the treated city decreased by 1.14% after the delivery speed improved. These results collectively suggest that faster delivery not only led to consumers increasing their purchases online, but also induced them to do so in smaller and more frequent orders. As a consequence, the platform would need to expend more resources to fulfill and deliver each item

it sells. Hence, with faster delivery, the platform would face not only higher total operational and environmental costs, but also higher costs *per* unit of revenue.

We explore and identify evidence for two mechanisms behind our main results. First, we propose the *order-splitting effect*, whereby as delivery speed improves, consumers may split their regular purchases on the platform into smaller and more frequent orders. This behavior is consistent with classic time-discounting preferences (Samuelson 1937, Frederick et al. 2002), which suggest that as delivery speed increases, the consumers’ marginal utility of getting an item sooner increases, and so they would become less willing to wait to consolidate their orders. We analyze consumers’ regular purchases by identifying the product categories that each consumer purchased in the six months prior to the study period (“regular categories”). We show that, indeed, consumers spent significantly less within their regular categories in each order after delivery became faster. Importantly, we show that this reduction is driven by fewer purchase quantities per order, while the average price of items purchased in these categories remain unchanged. Second, we propose and verify the *category expansion effect*, whereby consumers—after observing faster delivery—expand the product categories they purchase online and, notably, may prefer to purchase new categories with a different shopping frequency compared to the old categories.

Finally, to illustrate the impact of these behavioral changes as a result of faster delivery, we conduct two back-of-the-envelope calculations to highlight how the changes in consumer order frequency and basket size could affect the platform’s operational costs and carbon emissions. First, we estimate that the marginal fulfillment cost for each order is around 7.1 RMB for the platform. If Alibaba implemented the delivery speed change across China (with similar effects on consumer ordering behavior), it would result in an annual surge of over 7.5 billion deliveries from the platform’s existing consumers, which would translate into an additional fulfillment cost of 53.3 billion RMB (7.7 billion USD) per year. Moreover, approximately 13.9 billion RMB (2.0 billion USD) of this additional cost can be directly attributed to consumers’ reduced basket sizes. In other words, if the company made cost projections based on previous consumer basket sizes, then it would be short by 13.9 billion RMB (or 26.1% of the projected fulfillment costs). Second, we conduct similar calculations for the platform’s carbon emissions. Our calibration suggests that each additional order on the platform would lead to an additional 224g of carbon emissions due to packaging and transportation alone. Hence, a China-wide implementation of the delivery speed improvement (with similar effects on consumer ordering behavior) would translate into at least 1.7 million metric tons of additional carbon emissions per year (439,746 of which is due to reduced basket sizes). These carbon emissions would need 77.3 million trees to offset and correspond to an additional carbon cost of 16.4 million USD for the platform.

Our results make several important contributions. To the extent of our knowledge, we are the first to empirically demonstrate the impact of faster delivery on consumer order frequency and basket size. Our results suggest that faster delivery not only leads to increased sales, but also causes substantial increases in platforms’ operational and environmental costs due to additional deliveries, which platforms should take into account when planning investments to improve delivery speed. In addition, we propose two novel mechanisms that could drive our results and obtain evidence consistent with these mechanisms, thus contributing to the literature that seeks to understand consumer behavior in the presence of logistics changes. Furthermore, our results highlight that as e-commerce platforms increase their delivery speed, their costs *per* unit of revenue will increase, due to consumers’ decreasing basket sizes. Hence, for platforms that wish to keep their operational costs and carbon emissions in check, they should invest even more in strategies that reduce delivery-related costs/emissions, such as smarter packaging, better warehouse management, or electric transportation vehicles. Policies that directly induce consumers to consolidate orders (e.g., offering to deliver multiple orders on the same day or providing a discount for items that ship together) could also be effective, as well as algorithms that automate order consolidation at the warehouse level.

The remainder of this paper is organized as follows. Section 2 reviews the related literature and our contributions. Section 3 describes the empirical setting and data. In Section 4, we demonstrate the causal impact of faster delivery on consumer purchase behavior. Section 5 explores two potential mechanisms underlying the empirical results. In Section 6, we discuss the impact of consumers’ behavioral changes on the company’s operational costs and carbon footprint. We conclude and discuss future research directions in Section 7.

2. Literature Review

Our work is related to three streams of literature in operations management: e-commerce logistics, platform operations and sustainable operations.

2.1. E-commerce Logistics

First, our work contributes to a rapidly growing empirical research that explores the role of logistics in e-commerce. Several prior empirical studies have investigated how different delivery strategies impact consumer behavior and sales performance of the e-commerce companies. For example, in omnichannel retailing, offering consumers the choice of purchasing online and picking up in stores reduces online sales but increases in-store sales (Gallino and Moreno 2014), and it increases the retailer’s sales dispersion (Gallino et al. 2016). Wu et al. (2024) show that managing delivery service internally can boost consumer purchases and improve sales due to enhanced delivery service quality and consumers’ trust. Wagner et al. (2023) demonstrate that consolidating deliveries of multiple

orders can improve consumers' satisfaction and reduce product returns in an online fashion retail context. Cachon et al. (2018) analyze the relationship between the free shipping threshold and consumer purchases. In addition, some researchers have focused on how operational levers in last-mile delivery can enhance consumer experience and contribute to online sales, including optimizing the timing of track-package activities in last-mile logistics (Bray 2023), providing doorstep deliveries (Lu et al. 2023, Zhan et al. 2023), offering consumers the flexibility to pick up orders in pick-up stations (Bai et al. 2022), and using autonomous vehicles (Han et al. 2024).

Particularly relevant to our research is the literature that examines the impact of delivery speed on consumer purchases and sales of online platforms. A number of studies have empirically shown that faster delivery speed can increase consumers' purchases and boost sales in online retail. For example, Fisher et al. (2019) show that improving delivery speed increases the online sales of an apparel retailer. Deshpande and Pendem (2022) show that faster delivery speed positively impacts consumer ratings of logistic service, which in turn increases online retailers' sales. Lim et al. (2023) also demonstrate the positive effect of faster delivery on consumer demand, but show that the magnitude of this effect may differ across different channels. Conversely, Cui et al. (2020a) find that the removal of fast and reliable delivery service leads to a decline in sales. Similar effects of faster delivery on consumer demand have also been identified in other contexts, such as meal delivery platforms (Mao et al. 2019, Cui et al. 2023). Moreover, consumer demand on online platforms can be impacted by the discrepancy between the estimated and actual delivery times (Cui et al. 2020b) and sellers' commitment of delivery service performance (Luo et al. 2020).

The existing literature (as summarized above) tend to focus solely on examining the impact of logistics service on the overall consumer demand on online platforms. As a result, they only highlight the benefit of improved services in increasing the platforms' revenue, but do not reveal the platforms' operational and environmental costs for serving that increased demand. To the best of our knowledge, our work is the first to empirically analyze the causal effect of delivery speed improvement on the more granular aspects of consumer ordering behavior such as consumers' purchase frequency and basket size. In particular, we contribute to this literature by showing that faster delivery not only induces consumers to spend more, but also encourages them to place smaller and more frequent orders, which would result in additional operational and environmental costs per unit of sale. We further explore and identify two novel mechanisms that drive this phenomenon, which can help platforms identify appropriate strategies to counter the increased costs.

2.2. Platform Operations

Our work adds to the empirical literature that study how the operational practices of online platforms, particularly online retail businesses, could impact consumers' behavior and firms' operational

performance. A number of studies have empirically investigated how consumers' purchase behavior can be impacted by various operational levers of online platforms, such as pricing strategies (Fisher et al. 2018, Moon et al. 2018), promotions (Zhang et al. 2020, Cohen et al. 2022), information provision (Gallino and Moreno 2018, Cui et al. 2019, Hu and Wang 2022, Calvo et al. 2023), customer engagement (Wang et al. 2022), and design of artificial intelligence algorithms (Ferreira et al. 2016, Feldman et al. 2022, Khern-am nuai et al. 2024), with implications for platform strategy. Our work contributes to this literature by identifying the potential operational cost and environmental impact brought by changes in consumer ordering behavior when delivery speed improves, and highlighting the importance of parallel initiatives to manage fulfillment costs and/or carbon emissions.

2.3. Sustainable Operations

Our work also contributes to the literature that studies sustainability issues in operations management (for recent reviews, see Agrawal et al. 2018, Atasu et al. 2019), which has studied how company operations influence environmental outcomes such as waste (e.g., Plambeck and Wang 2009, Keskin et al. 2023) and carbon emissions (e.g., Cachon 2014), as well as social welfare (e.g., Levi et al. 2024). Our work is particularly related to the papers that study the relationship between companies' operational decisions and their carbon emissions. For instance, Cachon (2014) studies a retailer's optimal retail store density decision and the associated emissions due to the retailer's store replenishment transportation, consumers' transportation, and store energy usage. Belavina et al. (2016) compare two revenue models for the online retailing of groceries (per-order versus subscription) and calculates the carbon emissions due to both food waste and delivery of groceries. Belavina (2021) also studies the impact of grocery-store density on carbon emissions, but focus exclusively on those due to food waste. Other research has studied strategies to incentivize adoption of green technologies and practices, such as electrical vehicles (Avci et al. 2014, Cohen et al. 2016, Lim et al. 2015), solar panels (Kundu and Ramdas 2022, Kaps and Netessine 2023), and shared mobility for last-mile delivery (Qi et al. 2018). For example, Avci et al. (2014) model the adoption and use of electric vehicles in a novel battery exchange system and compares the resulting carbon emissions and oil dependence with those of a traditional system. Qi et al. (2018) study the impact of crowdsourced last-mile delivery and show that this system does not offer significant emissions reduction compared to traditional truck-based delivery systems due to the lower loading capacity of cars (relative to trucks).

We contribute to this stream of literature in two ways. First, despite the growing prominence of the e-commerce industry, there is little existing research on sustainability issues in e-commerce. Therefore, our research makes a step towards filling this gap by studying the carbon emissions

implications of consumer purchase behavior shifts as a result of faster e-commerce delivery. Second, we empirically show and quantify a new source of tension between demand expansion and its environmental costs—the rising environmental costs to fulfill each unit of sale as delivery speed improves, thus highlighting an area where carbon-reduction initiatives could become more urgent as e-commerce companies grow.

3. Empirical Setting and Data

In this section we introduce our empirical setting and the data used in our study. We first describe our empirical context in Section 3.1 and our data in Section 3.2. We then elaborate on our identification strategy and verify its internal validity in Section 3.3.

3.1. Empirical Setting

To investigate the impact of faster delivery on consumer purchase behavior, we collaborate with Alibaba Group, the largest e-commerce company in China. In the fiscal year of 2022, Alibaba Group served approximately 1.31 billion consumers globally and generated a total gross merchandise value (GMV) of 8,317 billion RMB (1,204 billion USD) across its consumer-facing businesses (Alibaba Group 2022). Similar to Amazon, Alibaba engages in both direct retail sales, where Alibaba sell their procured products directly to consumers like traditional retailers, and third-party sales, where third-party sellers offer their goods on Alibaba’s platform. In particular, our research employs data obtained from *Tmall Supermarket*, an online supermarket which Alibaba operates mostly through direct retail sales. Established in 2011, Tmall Supermarket provides a wide range of consumer goods, including daily necessities, fast-moving consumer goods, and other general merchandise, to a vast consumer base across China. For the 2022 fiscal year, Alibaba’s direct retail sales in China yielded a revenue of nearly 261 billion RMB (38 billion USD), with Tmall Supermarket being a primary contributor to this figure (Alibaba Group 2022).

Tmall Supermarket relies on Cainiao Network, Alibaba’s own logistics arm, to provide its delivery services. Tmall Supermarket employs a threshold-based free shipping policy: For orders within designated weight limits, those valued at 88 RMB (12.74 USD) or above are granted free shipping, while those below this threshold are subject to a shipping fee of 5 RMB (0.72 USD). Overweight orders will incur an additional charge of 1 RMB (0.14 USD) per kilogram of excess weight. However, our data indicates that overweight orders account for less than 1% of total orders. The shipping policy is displayed on the Tmall Supermarket home page, as depicted in Figure 1a. At the checkout counter, consumers can easily view the applicable shipping fee and estimated delivery time for their orders, as demonstrated in Figure 1b.

Tmall Supermarket has been dedicated to improving its delivery speed as a primary way to enhance consumer experience and maintain competitiveness since its establishment. A crucial strategy employed by Tmall Supermarket, and Alibaba Group in general, to expedite delivery has been




Figure 1 Tmall Supermarket Homepage and Checkout Page

building new warehouses within or near cities. Establishing a new warehouse close to a city allows Tmall Supermarket to ship orders for that city and its surrounding areas from a local warehouse rather than a more distant one, which significantly reduces transportation distances and hence delivery times. In 2015, Tmall Supermarket had only 8 warehouses in China, but this number rose to 17 warehouses in 2016 and further expanded to 28 warehouses in 2017. This substantial growth contributed significantly to improvements in delivery speed: the coverage of Tmall Supermarket's same-day delivery service expanded from a handful of the largest cities in 2015 to 11 cities in 2016 and 34 cities in 2019 (Orient Securities 2019).

In this study, we exploit the opening of a new Tmall Supermarket warehouse in Zhengzhou, the capital of Henan Province in China, as a natural experiment to causally estimate the effects of improved delivery speed on consumer purchase behavior. In August 2017, Tmall Supermarket opened a new warehouse in Zhengzhou (hereafter referred to as the “Zhengzhou warehouse”) which exclusively handles orders within Henan Province. The order fulfillment for Zhengzhou was then transferred from a warehouse outside Henan Province to this local facility. As a result, the delivery time for orders in Zhengzhou decreased from an average of approximately 2.5 days to less than 2 days following the warehouse's opening. Meanwhile, the order fulfillment process and associated delivery times for orders in Jinan, the capital of the neighboring Shandong Province, remained unchanged. Specifically, the average delivery time for orders in Jinan remained constant at around 1 day (28 hours) before and after the opening of the Zhengzhou warehouse.³ Therefore, it is possible to leverage the opening of the Zhengzhou warehouse as an exogenous shock to the delivery speed in Zhengzhou and causally identify the impact of delivery speed improvement on consumer purchase behavior by comparing consumer purchases in Zhengzhou (treated city) and Jinan (control city).

³ We also examined other provincial capitals in neighboring provinces of Zhengzhou as candidate controls. Jinan is the only city whose pre-intervention trend in delivery time does not differ significantly from that of Zhengzhou. Combined with Jinan's socioeconomic similarity to Zhengzhou (e.g., in population size and GDP (National Bureau of Statistics of China 2023)), we identified Jinan as the most appropriate control city for our study.

Our empirical setting offers several significant advantages that lend credibility to the causal inferences drawn from our natural experiment. First, the timings and locations of new warehouse openings followed a pre-determined plan for expanding Alibaba’s logistics network across the country. These decisions were made at the corporate logistics planning level in advance and did not depend on recent or anticipated local sales performance.⁴ Accordingly, the new warehouse openings were not triggered by local consumer demand. Second, the company does not announce the launch of new warehouses to consumers in advance, nor does it publicize such operational changes through marketing campaigns or press releases. The warehouse transition was executed entirely on the supply side, with the local orders rerouted from the previous facility to the new one as planned by the company. This prevents consumers from anticipating improvements in delivery speed and adjusting their ordering behaviors. These institutional features indicate that the opening of the Zhengzhou warehouse was exogenous with respect to the local consumer demand. Third, each city independently operates its logistics network and delivery services, without operational interference across cities. Therefore, the opening of the Zhengzhou warehouse was an exogenous shock to the delivery times of Tmall Supermarket orders in Zhengzhou but does not impact the deliveries of any orders in other cities in the neighboring provinces.

In addition, Tmall Supermarket employs a centralized operations model. Their key operational policies are determined centrally by the platform (including product assortment and delivery policies) or by the suppliers (including pricing and promotions) and implemented uniformly across the country, and thus these policies also remained identical in the treated and control cities throughout our study period. Moreover, each local warehouse is designed to carry the full (and identical) product assortment. Each consumer order is fulfilled by a single warehouse assigned by the company based on geographic location and is not split across multiple warehouses. The new warehouse opening thus did not change the set of products available to consumers in either the treated or the control city. Therefore, the opening of the Zhengzhou warehouse could impact consumer purchase behavior only through changes in delivery time, rather than through differences in other operational factors, such as product assortment, item pricing, or delivery charges between the two cities.

3.2. Data

Our primary dataset is the order-level data from Tmall Supermarket in Zhengzhou and Jinan between March 1, 2017 and February 28, 2018 (hereafter referred to as the “study period”). We restrict our sample to existing consumers who made at least one purchase during the 6 months before the study period (i.e., from September 1, 2016 to February 28, 2017; hereafter referred to as

⁴In fact, a primary determining factor in warehouse location and timing decisions was coordination between the company and local government authorities, which significantly impacts land acquisition approvals and permissions.

the “pre-study period”). For each order, we observe the consumer ID, zip code, items purchased, quantity of each item purchased, item prices, and the corresponding shipping fee. Note that we only focus on existing consumers since we would like to study consumers’ behavior changes with respect to delivery speed improvement rather than the overall impact on demand due to faster delivery. In other words, our estimates do not account for the potential effect of faster delivery on attracting more newcomers.

We divide the study period in our data sample into three periods, as illustrated in Figure 2: (1) the pre-intervention period from March 1, 2017 to July 31, 2017, covering five months before the Zhengzhou warehouse opening, (2) the intervention period from August 1, 2017 to September 30, 2017, corresponding to the period when the Zhengzhou warehouse was launched and order fulfillment was gradually transitioned from the previous facility, and (3) the post-intervention period, from October 1, 2017 to February 28, 2018, covering five months after the Zhengzhou warehouse opening, during which the transition was completed and the warehouse operated steadily to fulfill orders in Zhengzhou.⁵ To minimize potential bias in our estimation due to the order relocation and operational ramp-up during the new warehouse’s transition phase, we exclude the intervention period from our analysis and only use the data from the pre-intervention and post-intervention periods to estimate the treatment effects. Additionally, we exclude orders placed during major promotional campaigns on the platform, including the 618, Double 11, Double 12, and Chinese New Year Shopping Festivals, as consumer purchase behavior during these events may significantly differ from regular periods, and order deliveries could be substantially delayed due to the overwhelming order volume on the platform during these periods.⁶ Consequently, the dataset for our analysis comprises 946,079 orders from all existing consumers during the pre-intervention and post-intervention periods, including 456,804 orders from consumers in the treated city and 489,275 orders from consumers in the control city.

To examine consumer-level purchasing behavior in terms of overall spending and order frequency, we also transform the order-level data into a balanced panel at the consumer-month level (i.e., each observation corresponds to a given consumer in a specific month). We primarily examine three outcome variables in our analysis. At the consumer-month level, we measure consumer spending using the average daily expenditure in a month, $AvgDailyExpenditure_{im}$, computed as the total

⁵ At Alibaba, a warehouse opening involves a structured transition period rather than an instantaneous switch. Before the Zhengzhou warehouse went live, it underwent preparatory work (e.g., inventory stocking, system integration) and operational testing to ensure readiness. In late August, the warehouse began processing real consumer orders, which we mark as “warehouse opening.” The warehouse then initially operated with a relatively small order volume in a pilot mode. Over the subsequent weeks, order volume was progressively migrated from the previous warehouse to the new one and gradually scaled up. The complete migration of order flow was completed toward the end of October.

⁶ Our empirical findings remain qualitatively unchanged and statistically significant if these promotional periods are included in our analysis. See Table 8 in Appendix A for details.




Figure 2 Timeline of Our Data Sample

Table 1 Summary Statistics of Consumer Purchases in Zhengzhou (Treated City) and Jinan (Control City)

Variable	Granularity	Zhengzhou			Jinan		
		Observations	Mean	Standard deviation	Observations	Mean	Standard deviation
<i>AvgDailyExpenditure</i> ($\times 10^{-2}$)	consumer-month	4,035,140	44.36	179.59	2,962,290	66.59	238.44
<i>AvgDailyOrderCount</i> ($\times 10^{-2}$)	consumer-month	4,035,140	0.38	1.38	2,962,290	0.55	1.76
<i>BasketSize</i>	order	456,804	117.90	66.67	489,275	121.25	71.11

Note: This table reports the summary statistics of a consumer's average daily spending, average number of orders placed by a consumer per day, and the monetary value of an order during the pre-intervention and post-intervention periods for Zhengzhou and Jinan, respectively. For readability, *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100.

spending of consumer i in month m divided by the number of valid (i.e., non-promotional) days in month m . Similarly, we measure consumer order frequency using the average daily number of orders in a month, $AvgDailyOrderCount_{im}$, computed by the total number of orders placed by consumer i in month m divided by the number of valid days in month m .⁷ At the order level, we examine the basket size of an order, $BasketSize_{ijt}$, defined as the monetary value (in RMB) of an order j placed by consumer i on day t . Table 1 presents the summary statistics for the above three major outcome variables in our analysis in each city during the pre-intervention and post-intervention periods.

In addition, we acquired a separate dataset on delivery time for manipulation check, i.e., to verify the improvement in delivery speed in Zhengzhou following the opening of the Zhengzhou warehouse compared to Jinan. This dataset includes the average delivery time of all consumer orders from each district of Zhengzhou and Jinan for each month in both the pre-intervention and post-intervention periods.⁸ Specifically, there are 24 districts in the two cities (14 districts in Zhengzhou and 10 districts in Jinan) and 10 months in the pre-intervention and post-intervention periods (5 months in each of the pre-intervention and post-intervention periods). Therefore, we have $24 \times 10 = 240$ observations of average delivery time for each district-month pair.

⁷ We construct the outcome variables in this manner to ensure comparability of the outcome measurements across months with different calendar lengths, which enhances the precision of our empirical analysis. Our results remain robust if we use monthly expenditure and monthly order frequency instead; see Table 9 in Appendix A for details.

⁸ The delivery time data was available only at the district-month level at the time when we acquired it; more granular data was not accessible.

To summarize, our empirical analysis involves three levels of data granularity. We analyze consumer overall spending and order frequency using the balanced panel data constructed at the consumer-month level; basket size and other within-order characteristics using the order-level data; and the parallel trends tests as well as the manipulation check for delivery time using district-month level data specifically on delivery time.

3.3. Identification Strategy and Verification of Internal Validity

We exploit the opening of the Zhengzhou warehouse as an exogenous shock to the delivery speed in Zhengzhou and use the difference-in-differences (DID) approach to causally estimate the impact of faster delivery on consumer purchase behavior. To verify the internal validity of our identification strategy, we conduct two analyses: First, we examine the parallel trends assumption for the DID analysis, i.e., the treatment and control groups follow parallel trends in the outcome variables and delivery time during the pre-intervention period. Second, we perform a manipulation check to confirm that the delivery speed of Zhengzhou indeed improved significantly compared to that of Jinan after the Zhengzhou warehouse opened.

3.3.1. Verification of Parallel Trends Assumption We verify the parallel trends assumption on three variables: delivery time $DeliveryTime_{dm}$, defined as the average delivery time (in hours) of all orders placed by consumers from district d in month m , and the two consumer-month-level outcome variables on consumer expenditure and order frequency, $AvgDailyExpenditure_{im}$ and $AvgDailyOrderCount_{im}$, as defined in Section 3.2.⁹

We use three common approaches in the literature to verify the parallel trends assumption. First, following the literature in operations and marketing (Gallino et al. 2016, Fisher et al. 2019), we conduct a pre-trends test (i.e., test for pre-intervention differences in trends) and examine whether the differences in the trends of the dependent variables between the treated and control cities change linearly with time during the pre-intervention period. We use the following ordinary least squares (OLS) regression specifications for this test:

$$DeliveryTime_{dm} = \theta_0 + \theta_1 City_d \times Period_m + D_d + M_m + \epsilon_{dm}, \quad (1)$$

$$OutcomeVariable_{im} = \eta_0 + \eta_1 City_i \times Period_m + X_i + M_m + \epsilon_{im}, \quad (2)$$

where $OutcomeVariable_{im} \in \{AvgDailyExpenditure_{im}, AvgDailyOrderCount_{im}\}$. In Specifications (1) and (2), $Period_m \in \{1, 2, 3, 4, 5\}$ is a numerical variable that represents the chronological

⁹ Note that we do not conduct a parallel-trends test for basket size, because basket size is defined at the order level and thus forms an unbalanced panel—orders appear only conditional on purchase and each order is observed only once. This creates a selection issue. Therefore, a parallel-trends test for basket size does not provide a meaningful assessment of pre-intervention trends.

order of month m in the pre-intervention period (i.e., $Period_m = 1$ for March 2017 and $Period_m = 5$ for July 2017). $City_d$ (or $City_i$) is a binary variable that equals 1 if district d (or consumer i) is from the treated city and 0 otherwise. D_d , X_i , M_m represent district fixed effects, consumer fixed effects, and month fixed effects, respectively. ϵ_{dm} and ϵ_{im} are the error terms. Panel A of Table 6 in Appendix A presents the estimation results of Specifications (1) and (2) using the pre-intervention period data. We find that all the coefficient estimates of the interaction term $City \times Period$ are not statistically significant, suggesting no statistically significant differences in the trends of the dependent variables between the treatment and control groups during the pre-intervention period. This is consistent with the parallel trends assumption.

The limitation of the above test is the underlying linear assumption in the trends of the differences between the treatment and control groups. To overcome this limitation, we further conduct a placebo test (see, e.g., Gallino and Moreno 2014, Roth et al. 2023). Specifically, we introduce a “placebo treatment” during the pre-intervention period and estimate the effect of the placebo treatment on the dependent variables using the pre-intervention period data. If the parallel trends assumption holds, we should observe no significant effect of the placebo treatment. We set our placebo treatment to be on May 1, 2017, which is the middle of the pre-intervention period in our sample. Our results are qualitatively robust towards the selection of placebo treatment time. To conduct the placebo test, we use specifications that replace the variable $Period_m$ in Specifications (1) and (2) by the variable $PlaceboAfter_m$, which is a binary variable that equals 1 if month m is after the placebo treatment and 0 otherwise. The coefficients of the interaction term $City \times PlaceboAfter$ captures the effect of the placebo treatment. We present the placebo test results in Panel A of Table 6 in Appendix A. Similarly, we do not find any statistically significant effects of the placebo treatment on the dependent variables. This result is also consistent with the parallel trends assumption.

Finally, to visually present both the pre-intervention parallel trends as well as the dynamic treatment effects after the warehouse opening, we provide the event-study plots for delivery time, consumer expenditure and order frequency in Figure 3 in Appendix A, respectively. These plots show no statistically significant differential pre-trends between the treated and control cities. The figures thus provide further support for the parallel trends assumption and reinforce the validity of our DID identification strategy.

3.3.2. Manipulation Check for Delivery Speed Improvement We conduct a manipulation check to verify that the delivery speed of Zhengzhou indeed significantly improved compared to that of Jinan after the opening of the Zhengzhou warehouse. Figure 4 in Appendix A displays the average delivery time in the treated and control cities each month during the pre-intervention and

post-intervention periods. It demonstrates a noticeable reduction in the delivery time in Zhengzhou after the Zhengzhou warehouse opened, as compared to Jinan. We next employ the DID approach to causally estimate the change in delivery time in Zhengzhou followed by the local warehouse opening relative to Jinan. Specifically, we run the following OLS regression model:

$$DeliveryTime_{dm} = \gamma_0 + \gamma_1 City_d \times After_m + D_d + M_m + \epsilon_{dm}, \quad (3)$$

where $After_m$ is a binary variable that equals 1 if month m is during the post-intervention period and 0 otherwise. The definitions of the other variables are the same as those in Specification (1). The coefficient γ_1 of the interaction term $City \times After$ captures the impact of the Zhengzhou warehouse's opening on the delivery time in Zhengzhou. Panel B of Table 6 in Appendix A presents the estimation results from Specification (3). We find that the opening of the Zhengzhou warehouse significantly decreased the delivery time in Zhengzhou by 13.59 hours, equivalent to a relative decrease of 23.63% ($= 13.59/57.52$). This suggests that the opening of the Zhengzhou warehouse significantly improved the delivery speed in Zhengzhou, which confirms the validity of our identification strategy.

4. Empirical Results

In this section, we investigate how delivery speed improvement influences consumer purchase behavior on the platform. We first examine the impact of faster delivery on consumer spending in Section 4.1. We then show how faster delivery affects consumers' order frequency in Section 4.2. In Section 4.3, we analyze the effect of faster delivery on the basket size of consumer orders. Finally, we verify the robustness of our results using the synthetic difference-in-differences (SDID) approach and present several heterogeneous treatment effect analyses in Section 4.4.

4.1. Impact of Faster Delivery on Consumer Spending

We first causally estimate the effect of delivery speed improvement on consumers' spending using the DID approach. Specifically, we use the following OLS regression specification:

$$OutcomeVariable_{im} = \alpha_0 + \alpha_1 City_i \times After_m + X_i + M_m + \epsilon_{im}. \quad (4)$$

In Specification (4), $OutcomeVariable_{im}$ represents the outcome variable of interest and is $AvgDailyExpenditure_{im}$ as defined in Section 3.2 for our analysis in this subsection. The other variables are as defined in Section 3.3. The coefficient α_1 of the interaction term $City \times After$ captures the impact of the Zhengzhou warehouse opening and its resulting delivery speed improvement on consumers' expenditure. Column (1) in Panel A of Table 2 reports the results from Specification (4). We find that faster delivery facilitated by the Zhengzhou warehouse opening led to a significant

Table 2 Effects of Faster Delivery on Consumer Purchase Behavior

	Dependent variable		
	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$) (1)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$) (2)	$\log(\text{BasketSize})$ (3)
Panel A: DID results			
<i>City</i> $\times$ <i>After</i>	4.9872**** (0.3417)	0.0574**** (0.0026)	-0.0114**** (0.0023)
Relative effect size	11.05%	15.00%	-1.14%
Consumer fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-
Date fixed effects	-	-	Yes
Order control variables	-	-	Yes
Observations	6,997,430	6,997,430	946,079
Panel B: SDID results			
<i>City</i> $\times$ <i>After</i>	5.8542**** (1.1104)	0.0747**** (0.0084)	- -
Relative effect size	12.86%	19.46%	-
Consumer fixed effects	Yes	Yes	-
Month fixed effects	Yes	Yes	-
Observations	699,730	699,730	-

Note: This table reports the estimated effects of faster delivery on consumers' average daily expenditure (Column (1)), average daily number of orders (Column (2)), and basket size (Column (3)). Panel A reports the results of the DID approach and Panel B reports the results of the SDID approach. Relative effect sizes in Columns (1) and (2) are computed using the average pre-intervention outcomes of the treated city as the baseline, and the relative effect size in Column (3) in Panel A is computed by $(e^{-0.0114} - 1) = -0.0114$. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level in Panel A.

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

increase ($p < 0.001$) in the spending. Using the pre-intervention spending per treated consumer as the baseline, this increase translates to a relative growth of 11.05% in revenue for consumers in the treated city.

Our result demonstrates that, consistent with previous empirical studies (e.g., Fisher et al. 2019, Deshpande and Pendem 2022, Lim et al. 2023), faster delivery induces consumers to spend more on the online retail platform. Thus, this result confirms the effectiveness of faster delivery in boosting the platforms' revenue, which is a key rationale for the platforms' investment in improving delivery speed. However, this result does not answer the question of *how* consumers spend more on the platform—specifically, whether they place more order and if so, whether their order frequency increases faster or slower than expenditure. In the following subsection, we explore the effects of faster delivery on consumers' order frequency.

4.2. Impact of Faster Delivery on Consumer Order Frequency

We estimate how faster delivery affects consumer order frequency using Specification (4), with the outcome variable $OutcomeVariable_{im}$ being $AvgDailyOrderCount_{im}$ as defined in Section 3.2. As shown in Column (2) in Panel A of Table 2, the coefficient for *City* $\times$ *After* is positive and

statistically significant ($p < 0.001$). This implies that faster delivery resulted in a significant increase in the number of orders placed by a consumer per month, which amounts to a 15.00% increase in order frequency relative to the baseline frequency of the treated consumers in the pre-intervention period.¹⁰

This finding has several significant implications. First, it suggests that as delivery speed improves, the platform should plan for an increase in consumer orders, and the associated rise in fulfillment costs. In Section 6.1, we provide an analysis of how our estimated increase in order volume could translate into additional operational costs, using Alibaba’s current fulfillment process as a case study. Similarly, the increased volume of orders could potentially raise the platform’s environmental cost due to additional carbon emissions from packaging and transportation. We provide an estimate for this increase in Section 6.2.

Furthermore, we observe that the percentage increase in consumers’ purchase frequency (15.00%) is *greater* than that in their expenditure (11.05%). This comparison suggests that, with an increase in delivery speed, consumers may start ordering in smaller basket sizes. As discussed earlier, such behavior can undermine a platform’s profitability and exacerbate the logistical and sustainability challenges as it increases the packaging and shipping costs without generating proportionate revenue gains for the platform. We verify the existence of this behavior by analyzing the order-level data in the next subsection.

4.3. Impact of Faster Delivery on Consumer Basket Size

In this section, we employ order-level data to investigate whether consumers place smaller orders or not in response to improved delivery speed. We use the following specification for our analysis:

$$\log(BasketSize_{ijt}) = \beta_0 + \beta_1 City_i \times After_t + Z_j + X_i + T_t + \epsilon_{ijt}, \quad (5)$$

where $BasketSize_{ijt}$ is the monetary value (in RMB) of order j placed by consumer i in day t and $\log(\cdot)$ denotes natural logarithm. Z_j is a control variable indicating the period of day when an order is placed, which is classified into four intervals: morning (6am-12pm), afternoon (12pm-6pm), evening (6pm-12am) and midnight (12am-6am). X_i and T_t represent the consumer fixed effects and date fixed effects, respectively.

Column (3) in Panel A of Table 2 shows the estimation results from Specification (5). We find that following the opening of the Zhengzhou warehouse, basket size of consumers from the treated city significantly declined by 1.14% ($= e^{-0.0114} - 1$) ($p < 0.001$) relative to the control city. While

¹⁰ We also estimate the treatment effects on consumer spending and order frequency using the Poisson pseudo maximum likelihood (PPML) method as a robustness check, given the large proportions of zeros in these outcome variables. The estimated effects, presented in Table 7 in Appendix A, are statistically significant and qualitatively consistent with the OLS estimates.

being subject to selection effects, these results provide indirect evidence that as delivery speed improves, consumers tend to place smaller orders. Thus, although faster delivery may boost overall platform sales by encouraging more consumer purchases, it also induces smaller and more frequent orders from consumers.

4.4. Robustness of the Results and Heterogeneous Treatment Effect Analyses

In this section, we conduct additional analyses to examine the robustness of our empirical findings and heterogeneous treatment effects across several dimensions. We first validate our main empirical results using the synthetic difference-in-differences (SDID) approach in Section 4.4.1. We then investigate the heterogeneous treatment effects with respect to the magnitude of delivery speed improvement experienced by consumers (in Section 4.4.2) and consumer purchase frequency (in Section 4.4.3). Finally, we examine how the treatment effects evolve over the post-intervention period in Section 4.4.4.

4.4.1. Robustness Check Using the SDID Approach We examine the robustness of our previous empirical findings by employing the SDID approach (Arkhangelsky et al. 2021). As Roth (2022) has discussed, it is difficult to provide a direct test of parallel trends assumption in a semi-parametric form. Therefore, we use the SDID approach as a robustness test since this method allows for the relaxation of the parallel trends assumption in the DID method. Essentially, the SDID approach leverages pre-intervention data from control units to construct a synthetic control group for the average outcome of treated units by properly weighting each control unit and each pre-intervention period. As such, this approach obviates the need for the strict parallel trends assumption in the DID approach. We apply the SDID approach to our analyses in Sections 4.1 and 4.2, where our unit of observation is each consumer-month pair and we are interested in the effects of delivery speed improvement on the average daily expenditure, $AvgDailyExpenditure_{im}$, and average daily number of orders, $AvgDailyOrderCount_{im}$, by consumer i in month m , respectively. We note that this method cannot be applied to our third outcome variable $BasketSize_{ijt}$, because its unit of analysis is each order which cannot appear in both pre-intervention and post-intervention periods, and thus we cannot construct a synthetic control group from the data of the pre-intervention period and use it for estimation of the post-intervention period.

In SDID estimation, we first compute the unit weight $\hat{\omega}_i$ for each consumer $i \in \{1, \dots, N\}$ and time weight $\hat{\lambda}_m$ for each month $m \in \{1, \dots, M\}$, where N denotes the number of consumers and M denotes the number of months in our sample. Specifically, the unit weights $\hat{\omega}_i$'s are computed to align the pre-intervention trends in the weighted average outcome of the control consumers with those of the treated consumers. Similarly, the time weights $\hat{\lambda}_m$'s are computed so that the weighted average pre-intervention outcomes for the control consumers differ by up to a constant from their

average post-intervention outcomes. We refer the readers to Arkhangelsky et al. (2021) for more details about the methods to compute these weights. With the unit weights $\hat{\omega}_i$'s and time weights $\hat{\lambda}_m$'s computed, we then solve the following regression problem to estimate the treatment effect:

$$\{\hat{\tau}, \hat{\mu}, \hat{\phi}, \hat{\rho}\} = \arg \min_{\{\tau, \mu, \phi, \rho\}} \sum_{i=1}^N \sum_{m=1}^M (OutcomeVariable_{im} - \mu - \tau City_i \times After_m - \phi_i - \rho_m)^2 \hat{\omega}_i \hat{\lambda}_m, \quad (6)$$

where $OutcomeVariable_{im} \in \{AvgDailyExpenditure_{im}, AvgDailyOrderCount_{im}\}$, and ϕ_i and ρ_m capture consumer fixed effects and month fixed effects, respectively. The other variables are defined the same as in Section 3.3. The coefficient τ captures the impacts of faster delivery on the outcome variables. Due to computational limitations, we randomly sample 10% consumers in each city for our estimation. Our results remain qualitatively consistent across different random samples.

Panel B of Table 2 presents our estimation results from Specification (6). Columns (1) and (2) indicate that faster delivery significantly increased both consumers' expenditure and order frequency by 12.86% and 19.46%, respectively. Both results are consistent with the estimates in Sections 4.1 and 4.2 in terms of direction, magnitude and statistical significance. In addition, consistent with our earlier findings, the relative increase in order frequency is larger than that in consumer spending. These results confirm the robustness of our main findings in Sections 4.1 and 4.2, i.e., the delivery speed improvement due to the opening of the Zhengzhou warehouse significantly increased consumers' consumption online in terms of both spending and order frequency, and consumers may be ordering in smaller basket sizes as reflected in the higher increase in order frequency relative to spending after faster delivery.

4.4.2. Heterogeneous Treatment Effects by Magnitude of Delivery Speed Improvement In this section, we investigate the heterogeneous treatment effects of faster delivery across consumers experiencing different magnitudes of delivery speed improvements in their districts. Specifically, we split the consumers into two groups: those residing in districts experiencing low (bottom 50%) delivery speed improvements and those in districts with high (top 50%) delivery speed improvements. We then estimate the heterogeneous treatment effects of faster delivery across the two groups of consumers using the following specifications:

$$OutcomeVariable_{im} = \alpha_0 + \alpha_1 City_i \times After_m + \alpha_2 HighSpeedImprov_i \times After_m + X_i + M_m + \epsilon_{im}, \quad (7)$$

$$\log(BasketSize)_{ijt} = \beta_0 + \beta_1 City_i \times After_t + \beta_2 HighSpeedImprov_i \times After_t + Z_j + X_i + T_t + \epsilon_{ijt}, \quad (8)$$

where $HighSpeedImprov_i$ is a binary variable which equals 1 if consumer i resides in districts with high delivery speed improvements and 0 otherwise. The estimation results are presented in

Table 10 in Appendix A. We find that consumers in districts with high delivery speed improvement exhibited significantly greater increases in both expenditure and order frequency than those in districts with low improvement, while the reduction in basket size for the former group was smaller than the latter. This alignment between substantial service improvements and stronger consumer responses in their spending and order frequency supports our identification strategy that the new warehouse opening impacts the consumer purchase behavior through improved delivery speed.

We also estimate the effects of faster delivery on consumer expenditure, order frequency, and basket size for each group separately using Specifications (4) and (5) in the paper. Table 11 in Appendix A presents the estimation results. Consumers in high-improvement districts exhibited stronger responses: expenditure and order frequency increased by 12.11% and 15.61%, respectively, compared to 8.28% and 13.39% in low-improvement districts. Basket sizes decreased in both groups, dropping by 0.92% in high-improvement districts and 1.73% in low-improvement districts. These results further confirm our finding that larger delivery speed improvements induce stronger increases in consumer spending and order frequency.

4.4.3. Heterogeneous Treatment Effects by Consumer Order Frequency We next explore whether our observed treatment effects of faster delivery are present for both low- and high-frequency buyers, or are instead driven by one of the two groups. First, we examine how the treatment effects are moderated by consumers' purchase frequency, measured by the number of orders that a customer i placed during the pre-study period $PreOrderCount_i$, using the following OLS Specifications:

$$OutcomeVariable_{im} = \alpha_0 + \alpha_1 City_i \times After_m + \alpha_2 PreOrderCount_i \times After_m + \alpha_3 City_i \times After_m \times PreOrderCount_i + X_i + M_m + \epsilon_{im}, \quad (9)$$

$$\log(BasketSize)_{ijt} = \beta_0 + \beta_1 City_i \times After_t + \beta_2 PreOrderCount_i \times After_t + \beta_3 City_i \times After_t \times PreOrderCount_i + Z_j + X_i + T_t + \epsilon_{ijt}. \quad (10)$$

We estimate Specifications (9) and (10) using the consumer-month-level data and the order-level data, respectively, with $OutcomeVariable_{im} \in \{AvgDailyExpenditure_{im}, AvgDailyOrderCount_{im}\}$. The estimation results are reported in Table 12 in Appendix A. We find that the treatment effects in both total expenditure and order frequency are statistically significant across consumers with different purchase frequencies, and more frequent buyers exhibit larger increases in their expenditure and order frequency following the delivery speed improvement relative to less frequent buyers.¹¹ Meanwhile, the reduction in basket size due to faster delivery

¹¹ Note that the coefficient of $City \times After$ is interpreted as the treatment effect for consumers with no purchases during the pre-study period ($PreOrderCount_i = 0$), while our sample only contain consumers with at least 1 purchases during the pre-study period (i.e., $PreOrderCount_i \geq 1$). Therefore, we need to combine the coefficients of $City \times After$ and $City \times After \times PreOrderCount$ to obtain the baseline treatment effects for the consumers with minimum purchase frequency $PreOrderCount_i = 1$ in our sample.

is consistent across consumers with different purchase frequencies, and we do not observe statistically significant heterogeneity in this effect. These results suggest that the main treatment effects of faster delivery hold across consumers with different purchase frequencies and are even more pronounced for high-frequency buyers in terms of increased expenditure and purchase frequency.

In addition, we estimate the treatment effects separately for the two groups of low- and high-frequency buyers. We classify low-frequency buyers as those who made fewer than 3 purchases in the pre-study period (85% of the sample) and high-frequency buyers as those who made at least 3 purchases in the pre-study period (15% of the sample). The results, reported in Table 13 in Appendix A, show that for both groups of consumers, delivery speed improvement led to significant increases in total expenditure and order frequency and a significant reduction in average basket size. Consistent with our earlier findings, high-frequency buyers exhibit larger increases in their spending and order frequency than low-frequency buyers after faster delivery. Our results remain robust when we vary the cutoff used to classify the low- and high-frequency groups. These findings provide further evidence that our main empirical results—namely, that faster delivery induces consumers to purchase more through more frequent but smaller orders—hold across consumers with different purchase frequencies, with stronger spending and order frequency responses among high-frequency buyers.

4.4.4. Dynamics of the Treatment Effects Over the Post-intervention Period Finally, we analyze how the treatment effects evolve over the post-intervention period to examine whether the observed effects persist over time. We split the post-intervention period into two stages: an early stage (October-November 2017, the first two months after the intervention) and a late stage (December 2017-February 2018, the last three months of the post-intervention period). We then test whether the treatment effects in the late stage differ from those in the early stage by estimating the following specifications:

$$OutcomeVariable_{im} = \alpha_0 + \alpha_1 City_i \times After_m + \alpha_2 City_i \times LatePost_m + X_i + M_m + \epsilon_{im}, \quad (11)$$

$$\log(BasketSize)_{ijt} = \beta_0 + \beta_1 City_i \times After_t + \beta_2 City_i \times LatePost_t + Z_j + X_i + T_t + \epsilon_{ijt}, \quad (12)$$

where $LatePost_m$ (and $LatePost_t$) is a binary variable which equals 1 if month m (or date t) is in the late stage of the post-intervention period and 0 otherwise. The results are reported in Table 14 in Appendix A. We find that the increases in consumer expenditure and order frequency and the reduction in basket size already emerge during the early stage of the post-intervention period, and all these effects are statistically significant. Moreover, the increase in order frequency and reduction in basket size become even stronger in the late stage, while the increase in consumer expenditure does not change significantly. These results suggest that the effects of faster delivery on consumer purchase behavior are not just short-term impacts but persist over time, with consumers' tendency of placing more frequent and smaller orders even strengthening over time.

5. Possible Mechanisms

Our empirical results in Section 4 demonstrate that as delivery speed increases, consumers purchase more on the platform, but they also place orders more frequently and do so in smaller basket sizes. In particular, we note that while the fact that faster delivery can increase consumer demand on an online retail platform is well-established in the literature given consumers' demonstrated preference for fast delivery (Fisher et al. 2019, Deshpande and Pendem 2022, Lim et al. 2023), the questions of how that increased demand converts into changes in consumer ordering behaviors and what drives those behavioral changes have received scarce attention. In this section, we explore two mechanisms that could explain the observed purchasing pattern changes as delivery speed improves.

First, we propose that faster delivery may induce consumers to split their regular orders into more frequent ones due to classic time discounting preferences (Samuelson 1937, Frederick et al. 2002). To illustrate, consider a simple model in which a consumer decides between purchasing a product of value v on the platform now versus delaying the purchase to the next day (possibly consolidating the purchase with other demands that will arise). The consumer applies a discounting factor δ^t to the utility of a product received t days in the future, with $\delta \in (0, 1)$. Hence, the consumer's utility gap between purchasing the product now versus the next day is given by $v\delta^L - v\delta^{L+1}$, where L represents the delivery time. Notice that this utility gap is decreasing in L , which means that as delivery speed improves (i.e., as L decreases), the value of ordering an item one day sooner increases. Intuitively, this implies that a consumer may be willing to wait an additional day for a product that arrives (say) 30 days later, but would be much less willing to do so for a product that can be delivered tomorrow. Correspondingly, as delivery speed increases, consumers may become more "impatient" and less likely to wait to consolidate their orders, resulting in a higher purchase frequency for the same purchasing amount. We refer to this effect as the *order-splitting effect*.

Second, we propose that as consumers observe faster delivery on the platform, the *nature* of their demand may change, in the sense that consumers may start purchasing new categories of products online that they did not purchase before. In this case, the frequency and basket size with which they prefer to purchase these new categories may differ from those for existing categories. For instance, consumers who were accustomed to buying non-perishable goods every other Friday might begin purchasing smaller amounts of perishable items every Wednesday as delivery speed improves. In this case, the increase in consumers' order frequency is attributed not only to the accelerated purchasing rate of new categories but also to the desynchronization of ordering behaviors between existing and new categories. Therefore, faster delivery speed can lead to the overall increase in purchase frequency and reduction in basket size that we observe in our data as delivery speed increases. We refer to this effect as the *category expansion effect*.

In the next two subsections, we present supporting evidence for both mechanisms described above.

5.1. Order-Splitting Effect

We define the order-splitting effect as consumers splitting their regular orders into more frequent ones. To identify consumers' "regular orders", we separate each consumer's purchases into categories that they had purchased in the pre-study period ("regular categories") and those they had not ("new categories"). We then estimate how faster delivery impacts consumer purchase behavior in their regular categories (which we interpret as a proxy for their regular demand) at both the consumer-month and order level. Specifically, at the consumer-month level, for each consumer i and month m , we examine the average daily expenditure of consumer i on their regular categories in month m ($ExpenditureRC_{im}$) and average daily number of orders containing the regular categories of consumer i in month m ($OrderCountRC_{im}$). At the order level, we examine each order j placed by consumer i on day t that contains at least one regular category of consumer i and track, within these categories, (i) the total purchase amount ($AmountRC_{ijt}$), (ii) the number of items purchased ($ItemNumRC_{ijt}$), and (iii) the average item price ($ItemPriceRC_{ijt}$). We then estimate the following OLS regressions:

$$OutcomeVariable_{im} = \alpha_0 + \alpha_1 City_i \times After_m + X_i + M_m + \epsilon_{im}, \quad (13)$$

$$OutcomeVariable_{ijt} = \beta_0 + \beta_1 City_i \times After_t + Z_j + X_i + T_t + \epsilon_{ijt}, \quad (14)$$

where $OutcomeVariable_{im} \in \{ExpenditureRC_{im}, OrderCountRC_{im}\}$ and $OutcomeVariable_{ijt} \in \{\log(AmountRC_{ijt}), \log(ItemNumRC_{ijt}), \log(ItemPriceRC_{ijt})\}$. X_i , M_m , T_t and Z_j represent consumer fixed effects, month fixed effects, date fixed effects, and period-of-day indicators, respectively. Note that we estimate Specification (14) using only orders containing at least one regular category of the corresponding consumer.

The estimation results are presented in Tables 3 and 4. We find that, after delivery speed improvement, consumers' total expenditure and order frequency of regular categories increased by 12.6% and 15.5%, respectively, with the latter increasing at a higher percentage (Columns (1) and (2) of Table 3). At the order level, consumers spent significantly less (by 1.23%) within their regular categories in each order (Column (1) of Table 4). This reduction in spending within the regular categories is primarily driven by the *fewer purchase quantities* per order (Column (2) of Table 4), while the average price of items purchased in these categories remains unchanged (Column (3) of Table 4). These results suggest that as delivery speed improves, consumers are indeed de-consolidating their regular purchases—purchasing smaller *quantities* in each order, which contributes to the decline in average basket size.

We also perform similar analysis for new categories. Specifically, we construct analogous outcome variables for the new categories, including the average daily expenditure on new

Table 3 Effects of Faster Delivery on Consumer Purchases in Regular and New Categories

	Dependent variable			
	<i>ExpenditureRC</i> ($\times 10^{-2}$) (1)	<i>OrderCountRC</i> ($\times 10^{-2}$) (2)	<i>ExpenditureNC</i> ($\times 10^{-2}$) (3)	<i>OrderCountNC</i> ($\times 10^{-2}$) (4)
<i>City</i> $\times$ <i>After</i>	3.9239**** (0.2489)	0.0490**** (0.0022)	1.0633**** (0.0174)	0.0083**** (0.1871)
Relative effect size	12.55%	15.52%	7.66%	12.49%
Consumer fixed effects	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes
Observations	6,997,430	6,997,430	6,997,430	6,997,430

Note: This table shows the estimated treatment effects of faster delivery on the consumers' average daily expenditure on regular categories (Column (1)), average daily number of orders containing regular categories (Column (2)), average daily expenditure on new categories (Column (3)), and average daily number of orders containing new categories (Column (4)). The relative effect size is computed based on the average outcome value of the treated city during the pre-intervention period. For readability, the coefficient estimates and standard errors of all the outcome variables here are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level.

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

Table 4 Effects of Faster Delivery on Order-level Purchase Behavior in Regular and New Categories

	Dependent variable			
	$\log(\text{AmountRC})$ (1)	$\log(\text{ItemNumRC})$ (2)	$\log(\text{ItemPriceRC})$ (3)	$\log(\text{AmountNC})$ (4)
<i>City</i> $\times$ <i>After</i>	-0.0123** (0.0062)	-0.0111** (0.0044)	-0.0013 (0.0062)	0.0193 (0.0138)
Relative effect size	-1.23%	-1.10%	-	-
Consumer fixed effects	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes
Order control variables	Yes	Yes	Yes	Yes
Observations	764,596	764,596	764,596	517,812

Note: This table reports the estimated treatment effects of faster delivery on order-level purchase behavior, including purchase amount within regular categories (Column (1)), number of items purchased within regular categories (Column (2)), average item price within regular categories (Column (3)), and purchase amount within new categories (Column (4)). Standard errors are robust and clustered at the individual consumer level.

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

categories ($\text{ExpenditureNC}_{im}$), average daily number of orders containing new categories (OrderCountNC_{im}), and purchase amount within the new categories in an order (AmountNC_{ijt}). We then estimate Specifications (13) and (14) for these new-category outcomes, following the same empirical approach as in our analysis above for the regular categories (here we estimate Specification (14) using only orders that contain at least one new category of the corresponding consumer). We find that for new categories, while total expenditure and order frequency increased (with the latter outpacing the former; see Columns (3) and (4) of Table 3), there is no significant change in spending within new categories in each order that contains at least a new category (Column (4) of Table 4). Instead, it appears that new demand contributes to the reduction in basket size via *new orders*, as we will explain in the next subsection.

Table 5 Effects of Faster Delivery on Number of Categories Purchased

	Dependent variable			
	<i>MonthCateNum</i> (1)	<i>OrderCateNum</i> (2)	<i>HasNewCate</i> (3)	<i>OnlyHasNewCate</i> (4)
<i>City × After</i>	0.0414**** (0.0015)	0.0140** (0.0066)	0.0081*** (0.0026)	0.0110**** (0.0018)
Relative effect size	18.63%	0.60%	1.51%	6.35%
Consumer fixed effects	Yes	Yes	Yes	Yes
Month fixed effects	Yes	-	-	-
Date fixed effects	-	Yes	Yes	Yes
Order control variables	-	Yes	Yes	Yes
Observations	6,997,430	946,079	946,079	946,079

Note: This table shows the estimated treatment effects of faster delivery on the total number of distinct product categories consumers purchase per month (Column (1)), the number of distinct product categories consumer purchase per order (Column (2)), whether the order contains new categories (Column (3)), and whether the order only contains new categories (Column (4)). The relative effect size is computed based on the average outcome value of the treated city during the pre-intervention period. Standard errors are robust and clustered at the individual consumer level.

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

5.2. Category Expansion Effect

In this section, we demonstrate evidence of the category expansion effect. We first show that faster delivery indeed induced consumers to purchase products from more categories in a month. To this end, we estimate Specification (4) using *MonthlyCategoryNum_{im}* as the outcome variable, where *MonthlyCategoryNum_{im}* represents the total number of distinct product categories purchased by consumer i in month m . Column (1) of Table 5 presents the regression results. We find that faster delivery significantly increased the number of distinct product categories consumers purchase per month by 18.63%. If we restrict our analysis to months when consumers placed at least one order, then the number of distinct product categories consumers purchase per month still significantly increased by 2.75% (p -value < 0.001). This result suggests that as delivery speed improves, a portion of the increased consumer demand may stem from new categories of products that consumers have not purchased before, resulting in a broader range of product categories purchased each month.

Next, we show that the increased number of categories purchased contributes to the increase in order frequency. Specifically, we first conduct regression analysis using Specification (5) with the outcome variable being *OrderCategoryNum_{ijm}*, which represents the number of distinct product categories in order j placed by consumer i in month m . Column (2) of Table 5 presents the results. It shows that faster delivery significantly increased the number of distinct product categories purchased by consumers in each order by 0.60%. Comparing this increase with the much larger increase in the total number of product categories purchased per month (18.63%, or 2.75% conditioning on consumers making a purchase in that month), our data suggests that as consumer expand their purchases into new categories when delivery speed improves, they may be purchasing them in separate orders rather than consolidating them into existing regular orders.

Next, to more directly examine how consumers purchase to satisfy their new demand under the category expansion effect, we analyze consumers' purchases of their new categories (as identified in Section 5.1). Specifically, we examine whether the likelihood of an order containing a new category or an order containing only new categories increases after delivery speed improves. To do so, we estimate Specification (14) on all the orders with $OutcomeVariable_{ijt} \in \{HasNewCate_{ijt}, OnlyHasNewCate_{ijt}\}$, where $HasNewCate_{ijt}$ is 1 if order j placed by consumer i on day t contains their new categories and 0 otherwise, and $OnlyHasNewCate_{ijt}$ is 1 if order j placed by consumer i on day t contains only their new categories and no regular categories and 0 otherwise. We report the results in Columns (3) and (4) in Table 5, which further confirm our interpretation above: The share of orders containing new categories increased significantly after the delivery speed improved, and more interestingly, the share of orders containing *only* new categories increased significantly. Thus, the category expansion effect contributes to the increase in consumer order frequency mainly via *new* orders comprising entirely of new categories. Moreover, given that the orders consisting exclusively of new categories have significantly lower purchase amounts than those containing regular categories (115.70 RMB vs. 120.57 RMB; p -value < 0.001 in the two sample t-test), consumers' expanded purchases into new categories and the resulting increased prevalence of orders with exclusively new categories also contributes to the reduced average basket size.

In summary, these empirical evidences collectively suggest that the category expansion effect indeed exists and contributes to the increased order frequency and smaller basket size of consumers after the delivery speed improvement.

6. Fulfillment Costs and Environmental Implications of Faster Delivery for the Platform

In this section, we provide two back-of-the-envelope calculations to illustrate the impact of consumer order changes when delivery speed improves. Specifically, we first estimate the impact on fulfillment costs, and then present similar estimations on delivery-related carbon emissions. We use company-provided data when available, and supplement our calculations with public data. We highlight the key findings in the sections below, and relegate the details to Appendix B.

6.1. Fulfillment Costs

First, we consider how the increase in consumer order frequency would impact the company's fulfillment costs, which consist of warehousing, transportation and packaging costs related to deliveries. To do so, we first identify and estimate the average variable costs related to fulfilling each order, and then multiply the sum of the variable costs per order by the increase in order frequency.¹²

¹² Note that during our sample period (2017-2018), Alibaba did not offer platform-initiated order consolidation, so each order was processed independently from the next. Even today, automatic consolidation is not available for the vast majority of orders at Alibaba.

We note that this approach does not take into account any change in fixed costs or variable cost per order due to the actual implementation of the delivery speed improvement in our context (i.e., opening of the new warehouse). This is because companies can improve their delivery speed in many ways besides building new warehouses, such as re-routing, improving warehouse assortment, or better coordinating the outbound shipping process with warehouse operations through scheduling adjustments or information system upgrade, each with very different cost implications. Thus, by isolating the cost increase due solely to changes in consumer ordering behavior, our estimates can be generally applicable and likely provide a conservative estimate to multiple scenarios.

We estimate that the average variable cost of fulfilling an order is 7.1 RMB (1.0 USD); see details in Appendix B.1. Given an annual order volume of 49.7 billion at Alibaba Group (Citic Securities 2019), a 15.00% increase in consumer order frequency as a result of faster delivery—if implemented across the platform—would result in an annual surge of 7.5 billion ($= 49.7 \times 15.00\%$) billion deliveries from the platform’s existing consumers. The extra fulfillment costs from these additional deliveries would be around $7.5 \times 10^9 \times 7.1 = 53.3$ billion RMB (7.7 billion USD) per year. Importantly, approximately $(15.00\% - 11.05\%) \times 49.7 \times 10^9 \times 7.1 = 13.9$ billion RMB (2.0 billion USD) of these additional costs can be directly attributed to the reduction in consumers’ basket sizes. In other words, if the company made cost projections based on previous consumer basket sizes, then it would be short by 13.9 billion RMB (or 26.1% of the projected fulfillment costs).

6.2. Estimation of Carbon Footprint

We take a similar approach to estimate the impact of consumers’ ordering behavioral changes on the company’s carbon footprint. Specifically, we estimate that on average, each additional order on the platform would lead to an additional 224g of carbon emissions due to transportation and packaging alone; see details in Appendix B.2. As such, a China-wide implementation of the delivery speed improvement would result in $7.5 \times 10^9 \times (224/10^6) = 1.7$ million metric tons per year. As per the European Environmental Agency, a mature tree can absorb about 22 kg of carbon dioxide annually (European Environment Agency 2011). Hence, over 77.3 million ($= 1.7 \times 10^9/22$) trees would be needed to offset the additional carbon emissions from this increased order volume. Given that the 2023 annual average carbon trading price stands at 9.62 USD per metric ton in China (S&P Global Commodity Insights 2024), these additional carbon emissions correspond to a carbon cost of $9.62 \times 1.7 \times 10^6 = 16.4$ million USD for the platform.

Moreover, similar to our cost calculations earlier, $439,746 (= (15.00\% - 11.05\%) \times 49.7 \times 10^9 \times (224/10^6))$ metric tons of the additional carbon emissions from faster delivery can be directly attributed to the reduction in basket size. In fact, if the platform can instead induce consumers to increase their basket size by one third, it can save 3.2 million ($= 49.7 \times 10^9 \times (1 + 15.00\%) \times (1 -$

$1/(1 + 1/3)) \times (224/10^6)$) metric tons of carbon emissions each year. This illustrates the importance of consumer basket size in determining the platform’s carbon emissions.

We conclude by discussing a few caveats to the estimations above. First, it is important to note that our calculations are not meant to capture the carbon footprint from opening a new warehouse per se, but rather to highlight the environmental consequences from consumers’ downstream behavioral responses to faster delivery (which could be caused by a local warehouse opening). As such, we do not include the additional emissions from the actual construction and operations of the new warehouse, nor account for potential supply-side efficiency gains that may help reduce the carbon emissions per order (e.g., improved routing due to well-placed infrastructure, or assortment changes). Second, our carbon calculations are from the platform’s perspective,¹³ and hence do not reflect the entire lifecycle of the products sold (e.g., product manufacturing, third-party supplier logistics, and consumer use). A more comprehensive assessment may need to take into account how consumers use products, or the substitution between online and offline shopping behavior (see Walmart 2017, Carbonstop 2021 for several relevant analyses).

7. Conclusion

In this paper, we empirically investigate the causal effect of faster delivery on consumers’ purchase behavior. We utilize the opening of a local warehouse in Zhengzhou as an exogenous shock to consumers. Through our DID analysis on consumer-month-level data, we show that the half-day shipping time reduction brought by the opening of the new warehouse led to 11.05% more consumer spending and induced consumers to purchase 15.00% more frequently. Additional DID analysis on order-level data shows that average basket size decreased by 1.14%. Collectively, our empirical results show that although faster delivery can be an effective strategy to increase consumer spending on platforms, it also leads to more frequent and smaller orders. Our HTE analyses further suggest that faster delivery’s impact on spending and order frequency is greater for consumers experiencing greater delivery speed improvement or with higher purchase frequency. In addition, we find that the main effects persist over time, with consumers’ tendency of placing more frequent and smaller orders even strengthening over time. We identify and present evidence for two mechanisms driving our main results: order-splitting of regular purchases, and category expansion into new purchases. Finally, we provide estimates for the operational and environmental costs of these consumer behavioral changes in response to faster delivery.

Our findings have important implications for e-commerce platforms. We show that as these platforms invest in improving delivery speed, they should also plan for the substantial increase in consumer order frequency and the associated operational and environmental costs. Initiatives to

¹³ We follow a similar carbon accounting framework as used by Amazon (2021); see Appendix B.2.

directly reduce delivery-related costs and emissions, such as smarter packaging to improve the fill rate (Ahmad et al. 2022), improving fulfillment efficiency at warehouses, and switching to electric transportation vehicles, could become more important. In addition, platforms could consider policies to induce order consolidation. On the consumer side, platforms could offer incentives for greener delivery-time selection (e.g., offering to deliver multiple orders on the same day), provide a discount for items that ship together, or coordinate product promotions to encourage order schedule synchronization across different product categories. On the platform side, technology-driven solutions such as better predictive demand modeling, improved warehouse assortment management, and automated order consolidation could also help reduce the cost burdens associated with consumer behavioral changes in response to faster delivery.

In general, our work helps build an understanding of the causal effect of faster delivery on consumer purchase behavior, the underlying mechanisms, and the environmental and managerial implications. Our work has several limitations, which open up interesting avenues for future research. First, we do not have access to consumer characteristics in our data and therefore cannot identify whether there are heterogeneous treatment effects of faster delivery across different groups of consumers depending on their demographics. It would be interesting to see if such differences exist, and if so, how platforms can make use of this information to design more sophisticated policies that can achieve environmental sustainability without negatively impacting the platform's revenue. Second, due to data limitations, our work focuses on the impact of faster delivery on the behavior of the platform's existing consumers. Future research can further examine the behavioral tendencies of new consumers attracted by faster delivery, and whether they differ from those of existing consumers. Third, as we discussed at the end of Section 6.2, our carbon emission estimation represents one piece of a much more complex environmental puzzle, and more detailed field data and analyses are needed to fully understand the comprehensive environmental implication of faster delivery. Finally, it is worth exploring and empirically testing the effectiveness of various methods to nudge consumers to consolidate orders (such as Amazon Day Delivery), as well as consumers' acceptance of platform-initiated order consolidation.

References

- Agrawal VV, Atasu A, Van Wassenhove LN (2018) OM Forum — New opportunities for operations management research in sustainability. *Manufacturing & Service Operations Management* 21(1):1–12.
- Ahmad S, Utomo DS, Dadhich P, Greening P (2022) Packaging design, fill rate and road freight decarbonisation: A literature review and a future research agenda. *Cleaner Logistics and Supply Chain* 4:100066.
- Alibaba Group (2022) Fiscal Year 2022 Annual Report. <https://data.alibabagroup.com/ecms-files/886023430/c330302f-bfdd-4c79-a5ac-614446292e68.pdf>.
- Amazon (2021) Sustainability Report. <https://sustainability.aboutamazon.com/pdfBuilderDownload?name=amazon-sustainability-2020-report>.

- Amazon (2023) Reaching Net Zero Carbon by 2040. <https://sustainability.aboutamazon.com/carbon-methodology.pdf>.
- Arkhangelsky D, Athey S, Hirshberg DA, Imbens GW, Wager S (2021) Synthetic difference-in-differences. *American Economic Review* 111(12):4088–4118.
- Atasu A, Corbett CJ, Huang X, Toktay LB (2019) Sustainable operations management through the perspective of Manufacturing & Service Operations Management. *Manufacturing & Service Operations Management* 22(1):146–157.
- Avci B, Girotra K, Netessine S (2014) Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Science* 61(4):772–794.
- Bai B, Chan T, Zhang D, Zhang F, Chen Y, Hu H (2022) The value of logistic flexibility in e-commerce Working paper.
- Belavina E (2021) Grocery store density and food waste. *Manufacturing & Service Operations Management* 23(1):1–18.
- Belavina E, Girotra K, Kabra A (2016) Online grocery retail: Revenue models and environmental impact. *Management Science* 63(6):1781–1799.
- Blanco EE, Sheffi Y (2017) Green logistics. *Sustainable supply chains: a research-based textbook on operations and strategy*, chapter 7, 147–187.
- Bloomberg (2022) Amazon Plans to Invest \$1 Billion in Logistics and Robotics Companies. <https://www.bloomberg.com/news/articles/2022-04-21/amazon-to-invest-1-billion-in-logistics-robotics-companies?embedded-checkout=true>.
- Bray RL (2023) Operational transparency: Showing when work gets done. *Manufacturing & Service Operations Management* 25(3):812–826.
- Cachon GP (2014) Retail store density and the cost of greenhouse gas emissions. *Management Science* 60(8):1907–1925.
- Cachon GP, Gallino S, Xu J (2018) Free shipping is not free: A data-driven model to design free-shipping threshold policies. *Available at SSRN 3250971*.
- Calvo E, Cui R, Wagner L (2023) Disclosing product availability in online retail. *Manufacturing & Service Operations Management* 25(2):427–447.
- Carbonstop (2021) The GHG Emissions of Shopping: E-commerce Vs. Bricks and Mortar in China. <https://www.efchina.org/Reports-en/ghg-emissions-of-shopping-e-commerce-vs-bricks-and-mortar-in-china>.
- Citic Securities (2019) Comparative Analysis of Pinduoduo and Taobao in Five Dimensions. https://www.sohu.com/a/348274476_100020617.
- Cohen MC, Fiszer MD, Kim BJ (2022) Frustration-based promotions: Field experiments in ride-sharing. *Management Science* 68(4):2432–2464.
- Cohen MC, Lobel R, Perakis G (2016) The impact of demand uncertainty on consumer subsidies for green technology adoption. *Management Science* 62(5):1235–1258.
- Cui R, Li M, Li Q (2020a) Value of high-quality logistics: Evidence from a clash between SF Express and Alibaba. *Management Science* 66(9):3879–3902.
- Cui R, Lu Z, Sun T, Golden JM (2020b) Sooner or later? Promising delivery speed in online retail. *Manufacturing & Service Operations Management* 26(1).
- Cui R, Zhang DJ, Bassamboo A (2019) Learning from inventory availability information: Evidence from field experiments on amazon. *Management Science* 65(3):1216–1235.
- Cui R, Zhang W, Zheng Z (2023) Market thickness and delivery speed in food-delivery platforms. *Kelley School of Business Research Paper 2023-4419279*.
- Deshpande V, Pendem PK (2022) Logistics performance, ratings, and its impact on customer purchasing behavior and sales in e-commerce platforms. *Manufacturing & Service Operations Management* 25(3):827–845.

- Essense Securities (2021) Logistics Cost Report. <https://23813463.s21i.faiusr.com/61/ABUIABA9GAAGqNP1gwYo80jnigI.pdf>.
- European Environment Agency (2011) Forests, health and climate change. <https://www.eea.europa.eu/articles/forests-health-and-climate-change#:~:text=Trees%20also%20help%20tackle%20climate,of%20pollutants%20from%20the%20air>.
- Feldman J, Zhang DJ, Liu X, Zhang N (2022) Customer choice models vs. machine learning: Finding optimal product displays on alibaba. *Operations Research* 70(1):309–328.
- Fernández Briseño D, Chegut A, Glennon E, Scott J, Yang J (2020) Retail carbon footprints: Measuring impacts from real estate and technology. https://realestateinnovationlab.mit.edu/research_article/retail-carbon-footprints-measuring-impacts-from-real-estate-and-technology/.
- Ferreira KJ, Lee BHA, Simchi-Levi D (2016) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & service operations management* 18(1):69–88.
- Fisher M, Gallino S, Li J (2018) Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management science* 64(6):2496–2514.
- Fisher ML, Gallino S, Xu JJ (2019) The value of rapid delivery in omnichannel retailing. *Journal of Marketing Research* 56(5):732–748.
- Frederick S, Loewenstein G, O’donoghue T (2002) Time discounting and time preference: A critical review. *Journal of Economic Literature* 40(2):351–401.
- Gallino S, Moreno A (2014) Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science* 60(6):1434–1451.
- Gallino S, Moreno A (2018) The value of fit information in online retail: Evidence from a randomized field experiment. *Manufacturing & Service Operations Management* 20(4):767–787.
- Gallino S, Moreno A, Stamatopoulos I (2016) Channel integration, sales dispersion, and inventory management. *Management Science* 63(9):2813–2831.
- Greenhouse Gas Protocol (2004) The greenhouse gas protocol: A corporate accounting and reporting standard. *World Resources Institute and World Business Council for Sustainable Development* .
- Guardian (2020) Amazon profits surge as investment in faster shipping pays off. <https://www.theguardian.com/technology/2020/jan/30/amazon-profits-surge-investment-faster-shipping>.
- Han BR, Li M, Zhang Y, Li P (2024) Value of autonomous last-mile delivery: Evidence from alibaba. *Available at SSRN 4730263* .
- Hu K, Wang YI (2022) More or less: How information richness affects our decision consistency. *Vanderbilt Owen Graduate School of Management Research Paper* .
- Kaps C, Netessine S (2023) Residential battery storage - Reshaping the way we do electricity. *Available at SSRN 4219247* .
- Keskin NB, Li C, Song JSJ (2023) The blockchain newsvendor: Value of freshness transparency and smart contracts. *Available at SSRN 3915358* .
- Khern-am nuai W, So H, Cohen MC, Adulyasak Y (2024) Selecting cover images for restaurant reviews: Ai vs. wisdom of the crowd. *Manufacturing & Service Operations Management* 26(1):330–349.
- Kundu A, Ramdas K (2022) Timely after-sales service and technology adoption: Evidence from the off-grid solar market in uganda. *Manufacturing & Service Operations Management* 24(3):1329–1348.
- Levi R, Rajan M, Singhvi S, Zheng Y (2024) Improving farmers’ income on online agri-platforms: Evidence from the field. *Manufacturing & Service Operations Management* .
- Lim MK, Mak HY, Rong Y (2015) Toward mass adoption of electric vehicles: Impact of the range and resale anxieties. *Manufacturing & Service Operations Management* 17(1):101–119.
- Lim SFW, Gao F, Tan TF (2023) Channel changes choice: An empirical study about omnichannel demand sensitivity to fulfillment lead time. *Management Science* .
- Lu Z, Cui R, Sun T, Wu L (2023) The value of last-mile delivery in online retail. *Available at SSRN 4590356* .

- Luo J, Rong Y, Zheng H (2020) Impacts of logistics information on sales: Evidence from Alibaba. *Naval Research Logistics (NRL)* 67(8):646–669.
- Mao W, Ming L, Rong Y, Tang CS, Zheng H (2019) Faster deliveries and smarter order assignments for an on-demand meal delivery platform. *Available at SSRN 3469015*.
- Moon K, Bimpikis K, Mendelson H (2018) Randomized markdowns and online monitoring. *Management Science* 64(3):1271–1290.
- National Bureau of Statistics of China (2023) National data. <https://data.stats.gov.cn/english/easyquery.htm?cn=E0105>.
- Oliver Wyman (2023) Delivery Decarbonization Pathway. <https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2023/may/Delivery%20Decarbonization%20Pathway-Long%20Report.pdf>.
- Orient Securities (2019) E-commerce fast moving consumer goods strategy analysis: The trend towards online retail is clear, and new retail consolidates while moving forward. http://pdf.dfcfw.com/pdf/H3_AP201908141343638000_1.pdf.
- Plambeck E, Wang Q (2009) Effects of e-waste regulation on new product introduction. *Management Science* 55(3):333–347.
- Qi W, Li L, Liu S, Shen ZJM (2018) Shared mobility for last-mile delivery: Design, operational prescriptions, and environmental impact. *Manufacturing & Service Operations Management* 20(4):737–751.
- Roth J (2022) Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights* 4(3):305–322.
- Roth J, Sant’Anna PH, Bilinski A, Poe J (2023) What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics* 235(2):2218–2244.
- Samuelson PA (1937) A note on measurement of utility. *The Review of Economic Studies* 4(2):155–161.
- S&P Global Commodity Insights (2024) Commodities 2024: China’s domestic carbon market set for revamp; Article 6 in limbo. <https://www.spglobal.com/commodityinsights/en/market-insights/latest-news/energy-transition/011724-chinas-domestic-carbon-market-set-for-revamp-in-2024-article-6-in-limbo>.
- US Census Bureau (2023) Quarterly Retail E-commerce Sales (4th Quarter 2022). <https://www2.census.gov/retail/releases/historical/ecomm/22q4.pdf>.
- US EPA (2020) Corporate GHG Inventorying and Target Setting Self-Assessment: V1.0. <https://www.epa.gov/sites/default/files/2020-04/documents/corporate-ghg-inventorying-and-target-setting-self-assessment.pdf>.
- Wagner L, Calvo E, Amorim P (2023) Better together! The consumer implications of delivery consolidation. *Manufacturing & Service Operations Management* 25(3):903–920.
- Walmart (2017) The emissions implications of modern retailing: Omnichannel vs. stores and online pure-plays. <https://cdn.corporate.walmart.com/00/5a/3c20743a4f0db2d00c452aebea95/omni-channel-emissions-modeling-whitepaperfinal04182017.pdf>.
- Wang Y, Lu LX, Shi P (2022) Does customer email engagement improve profitability? Evidence from a field experiment in subscription service retailing. *Manufacturing & Service Operations Management* 24(5):2703–2721.
- Wu B, Chen Y, Naik PA (2024) Express: How own delivery services influence customer behavior and sales in online retail? Building trust and improving delivery quality in digital economy. *Journal of Marketing*.
- Zhan Y, Zhang Z, Ge J (2023) Value of exclusive doorstep delivery in the last-100-meter distribution. *Available at SSRN 4352455*.
- Zhang DJ, Dai H, Dong L, Qi F, Zhang N, Liu X, Liu Z, Yang J (2020) The long-term and spillover effects of price promotions on retailing platforms: Evidence from a large randomized experiment on alibaba. *Management Science* 66(6):2589–2609.

Appendix A: Figures and Tables

Table 6 Verification of Parallel Trends Assumption and Manipulation Check

	Dependent variable		
	<i>DeliveryTime</i>	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$)
	(1)	(2)	(3)
Panel A: Verification of Parallel Trends Assumption			
Pre-trends test: <i>City</i> $\times$ <i>Period</i>	0.2603 (0.3596)	-0.2164 (0.1453)	0.0016 (0.0011)
Placebo test: <i>City</i> $\times$ <i>PlaceboAfter</i>	1.3482 (1.0623)	-0.4563 (0.4118)	0.0043 (0.0031)
District fixed effects	Yes	-	-
Consumer fixed effects	-	Yes	Yes
Month fixed effects	Yes	Yes	Yes
Observations	120	3,498,715	3,498,715
Panel B: Manipulation Check			
<i>City</i> $\times$ <i>After</i>	-13.5932**** (1.6922)	-	-
Relative effect size	23.63%	-	-
District fixed effects	Yes	-	-
Month fixed effects	Yes	-	-
Observations	240	-	-


Note: This table reports the results of parallel trends assumption tests and manipulation check. Panel A presents the results of verifying parallel trends assumption on delivery time (Column (1)), consumers' average daily expenditure (Column (2)) and average daily number of orders (Column (3)) using the pre-trends test and placebo test. Panel B presents the results of manipulation check on delivery time (Column (1)). For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the district level for delivery time (Column (1)) and clustered at the individual consumer level for the other two variables (Columns (2)-(3)).

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


Table 7 Effects of Faster Delivery on Consumer Purchases Using Poisson Pseudo Maximum Likelihood

	Dependent variable	
	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$)
	(1)	(2)
<i>City</i> $\times$ <i>After</i>	6.3266**** (0.6046)	9.3930**** (0.5478)
Relative effect size	6.53%	9.85%
Consumer fixed effects	Yes	Yes
Month fixed effects	Yes	Yes
Observations	6,997,430	6,997,430


Note: This table reports the estimated effects of faster delivery on consumers' average daily expenditure (Column (1)) and average daily number of orders (Column (2)) using the Poisson Pseudo Maximum Likelihood approach. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level.



(a) Event-Study Plot of Delivery Time



(b) Event-Study Plot of Consumer Average Daily Expenditure



(c) Event-Study Plot of Consumer Average Daily Order Count

Figure 3 Event-Study Plots




Figure 4 Average Delivery Time in Zhengzhou and Jinan Before and After the Zhengzhou Warehouse Opened

Table 8 Effects of Faster Delivery on Consumer Purchase Behavior (Including Purchases During Promotions)

	Dependent variable		
	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$) (1)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$) (2)	$\log(BasketSize)$ (3)
<i>City</i> $\times$ <i>After</i>	4.2951**** (0.3707)	0.0541**** (0.0026)	-0.0125**** (0.0023)
Relative effect size	9.51%	14.15%	-1.24%
Consumer fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-
Date fixed effects	-	-	Yes
Order control variables	-	-	Yes
Observations	6,997,430	6,997,430	1,014,239

Note: This table reports the estimated effects of faster delivery on consumers' average daily expenditure (Column (1)), average daily number of orders (Column (2)), and basket size (Column (3)), with the orders placed during the major promotional campaigns included in the analysis. Relative effect sizes in Columns (1) and (2) are computed using the average pre-intervention outcomes of the treated city as the baseline, and the relative effect size in Column (3) is computed by $(e^{-0.0125} - 1) = -0.0124$. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level.

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

Table 9 Effects of Faster Delivery on Consumer Monthly Purchase Behavior

	Dependent variable	
	<i>MonthlyExpenditure</i> (1)	<i>MonthlyOrderCount</i> (2)
<i>City × After</i>	0.0499**** (0.0034)	0.01853**** (0.0008)
Relative effect size	11.05%	15.92%
Consumer fixed effects	Yes	Yes
Month fixed effects	Yes	Yes
Observations	6,997,430	6,997,430

Note: This table reports the estimated effects of faster delivery on consumers' monthly spending (Column (1)) and monthly number of orders (Column (2)). Relative effect sizes are computed using the average pre-intervention outcomes of the treated city as the baseline. Standard errors are robust and clustered at the individual consumer level.

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

Table 10 Heterogeneous Treatment Effects of Faster Delivery by Magnitude of Delivery Speed Improvement

	Dependent variable		
	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$) (1)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$) (2)	$\log(\text{BasketSize})$ (3)
<i>City × After</i>	3.4016**** (0.4295)	0.0463**** (0.0032)	−0.0176*** (0.0036)
<i>HighSpeedImprov × After</i>	2.2763**** (0.3942)	0.0159**** (0.0030)	0.0083** (0.0039)
Consumer fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-
Date fixed effects	-	-	Yes
Order control variables	-	-	Yes
Observations	6,997,370	6,997,370	946,073

Note: This table reports the heterogeneous treatment effects of faster delivery on consumers' average daily expenditure (Column (1)), average daily number of orders (Column (2)), and basket size (Column (3)), moderated by the magnitude of delivery speed improvement experienced by consumers. Observations with unspecified or incorrectly specified district codes, for which the corresponding delivery speed improvement cannot be determined, are excluded from the analysis. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level.

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

Table 11 Heterogeneous Treatment Effects for Consumers Experiencing Different Magnitudes of Delivery Speed Improvement

	Group 1: Low Delivery Speed Improvement			Group 2: High Delivery Speed Improvement		
	<i>AvgDailyExpenditure</i>	<i>AvgDailyOrderCount</i>	$\log(BasketSize)$	<i>AvgDailyExpenditure</i>	<i>AvgDailyOrderCount</i>	$\log(BasketSize)$
	($\times 10^{-2}$)	($\times 10^{-2}$)		($\times 10^{-2}$)	($\times 10^{-2}$)	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>City</i> $\times$ <i>After</i>	3.4016**** (0.4295)	0.0463**** (0.0032)	-0.0173**** (0.0036)	5.6779**** (0.3666)	0.0622**** (0.0028)	-0.0092*** (0.0025)
Relative effect size	8.28%	13.39%	-1.72%	12.11%	15.61%	-0.92%
Consumer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-	Yes	Yes	-
Date fixed effects	-	-	Yes	-	-	Yes
Order control variables	-	-	Yes	-	-	Yes
Observations	4,186,840	4,186,840	612,385	5,772,820	5,772,820	822,963

Note: This table reports the heterogeneous treatment effects of faster delivery on the purchase behavior of consumers experiencing different magnitudes of delivery speed improvement. Columns (1)-(3) report the estimated effects of faster delivery on consumers' average daily expenditure, average daily number of orders, and basket size, respectively, for the consumers residing in districts with low delivery speed improvement. Columns (4)-(6) report results of similar analyses for the consumers residing in districts with high delivery speed improvement. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the consumer level.

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

Table 12 Heterogeneous Treatment Effects of Faster Delivery by Consumer Order Frequency During the Pre-Study Period

	Dependent variable		
	<i>AvgDailyExpenditure</i>	<i>AvgDailyOrderCount</i>	$\log(BasketSize)$
	($\times 10^{-2}$)	($\times 10^{-2}$)	
	(1)	(2)	(3)
<i>City</i> $\times$ <i>After</i>	-0.1642 (0.8613)	0.0016 (0.0069)	-0.0115**** (0.0032)
<i>After</i> $\times$ <i>PreOrderCount</i>	-8.5225**** (0.3652)	-0.0779**** (0.0029)	0.0009*** (0.0003)
<i>City</i> $\times$ <i>After</i> $\times$ <i>PreOrderCount</i>	1.5069*** (0.5304)	0.0192**** (0.0042)	0.0004 (0.0006)
Consumer fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-
Date fixed effects	-	-	Yes
Order control variables	-	-	Yes
Observations	6,997,430	6,997,430	946,079

Note: This table reports the heterogeneous treatment effects of faster delivery on consumers' average daily expenditure (Column (1)), average daily number of orders (Column (2)), and basket size (Column (3)), moderated by consumers' order frequency during the pre-study period. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level.

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

Table 13 Effects of Faster Delivery for Low- and High-frequency Buyers

	Group 1: Low-frequency Buyers			Group 2: High-frequency Buyers		
	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$) (1)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$) (2)	$\log(BasketSize)$ (3)	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$) (4)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$) (5)	$\log(BasketSize)$ (6)
<i>City</i> $\times$ <i>After</i>	1.8214**** (0.3015)	0.0266**** (0.0023)	-0.0099*** (0.0033)	11.0731**** (1.3912)	0.1263**** (0.0104)	-0.0106*** (0.0033)
Relative effect size	5.67%	9.56%	-0.98%	8.22%	11.48%	-1.06%
Consumer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-	Yes	Yes	-
Date fixed effects	-	-	Yes	-	-	Yes
Order control variables	-	-	Yes	-	-	Yes
Observations	5,927,190	5,927,190	556,689	1,070,240	1,070,240	389,390

Note: This table reports the heterogeneous treatment effects of faster delivery for the low- and high-frequency buyers separately. Columns (1)-(3) report the estimated effects of faster delivery on consumers' average daily expenditure, average daily number of orders, and basket size, respectively, for the low-frequency buyers. Columns (4)-(6) report results of similar analyses for the high-frequency buyers. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the consumer level.

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

Table 14 Dynamics of the Effects of Faster Delivery Over the Post-Intervention Period

	Dependent variable		
	<i>AvgDailyExpenditure</i> ($\times 10^{-2}$) (1)	<i>AvgDailyOrderCount</i> ($\times 10^{-2}$) (2)	$\log(BasketSize)$ (3)
<i>City</i> $\times$ <i>After</i>	5.0276**** (0.4104)	0.0464**** (0.0032)	-0.0072*** (0.0027)
<i>City</i> $\times$ <i>LatePost</i>	-0.0672 (0.4326)	0.0182**** (0.0042)	-0.0080** (0.0032)
Consumer fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	-
Date fixed effects	-	-	Yes
Order control variables	-	-	Yes
Observations	6,997,430	6,997,430	946,079

Note: This table reports the dynamics of the treatment effects of faster delivery on consumers' average daily expenditure (Column (1)), average daily number of orders (Column (2)), and basket size (Column (3)) over the post-intervention period, by comparing the early (first two months) and late (last three months) stages. For readability, the coefficient estimates and standard errors of *AvgDailyExpenditure* and *AvgDailyOrderCount* are rescaled by multiplying the original values by 100. Standard errors are robust and clustered at the individual consumer level.

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

Table 15 Sources of Fulfillment Costs and Carbon Emissions in Online Shopping

Sources				Cost and Emission Type	Variable Fulfillment Cost (RMB per order)	Variable Carbon Emission (gram per order)
Order Processing	Facilities	Platform offices		Fixed	-	-
	Information processing	Data center		Fixed	-	-
Order Fulfillment	Warehousing	Facilities and assets		Fixed	-	-
		Labor		Variable	1.5	-
	Transportation	Line haul	Labor and vehicles	Fixed	-	-
			Fuel	Variable	0.4	74.2
		Last-mile delivery	Vehicles	Fixed	-	-
			Labor	Variable	2 – 4.9	0
			Electricity	Variable	< 0.01	4.5
	Packaging	Packaging material		Variable	1.5 – 2.0	145.3

Appendix B: Fulfillment Costs and Carbon Emission Calibrations

In this appendix, we provide detailed calibrations of the average fulfillment cost and carbon emissions the company would incur for each additional order. We do so by systematically breaking down all relevant cost and emission elements into fixed and variable components, and discussing our estimates for each of the variable components.

B.1. Fulfillment Costs

As shown in Table 15, we consider the fulfillment costs from each of the three phases of an order: (i) warehousing, (ii) transportation, and (iii) packaging.

Warehousing: Fulfillment costs in warehousing are those associated with operating warehouses and processing orders within the warehouses. A major contributor to these costs is the facilities and assets (e.g., warehouse rental, usage and depreciation of assets, power consumption, etc.). Another key factor is labor, the cost of which typically scales with the volume of orders handled by employees. According to Alibaba, an increase of 15.00% in order volume does not lead to expansion of warehouses or a substantial increase in asset costs of the warehouses. Therefore, we classify the costs associated with facilities and assets as fixed, while the labor cost is considered a variable cost.

Specifically, the average warehouse labor cost is estimated to be 1.5 RMB per order at JD.com, which is another major e-commerce company in China (Essense Securities 2021). Given that JD.com and Alibaba Group operate under similar circumstances (e.g., labor wage, productivity), we adopt the same estimate here.

Transportation: The fulfillment cost for shipping orders from warehouses to consumers comprise two primary components: line haul (i.e., transportation from warehouses to the last-mile delivery station), and last-mile delivery (i.e., transportation from the delivery station to consumer homes). Line haul costs are driven by factors such as labor, vehicles and fuel. We learn from Alibaba Group that the current vehicle fill

rate averages at 75% for Tmall Supermarket orders. Accordingly, an increase of 15.00% in order volume will not significantly affect the labor, vehicles and delivery schedules in line haul operations. Thus, we treat the costs associated with labor and vehicles in line haul as fixed. However, costs related to fuel in line haul are classified as variable as additional orders increase weights carried and hence the overall fuel consumption. The fuel cost during line haul of JD.com is estimated to be 0.4 RMB on average (Essense Securities 2021), which we adopt here.

Lastly, costs in last-mile delivery may come from labor, vehicles (mostly electric mopeds), and electricity consumed during delivery. According to Alibaba, couriers' salaries primarily depend on the number of orders delivered, and a 15.00% surge in volume would imply more delivery trips, though not necessarily a large fleet. Consequently, we categorize the labor and electricity costs as variable, and vehicle costs as fixed. According to a local Alibaba courier job post, the piecewise rate in last-mile delivery is 2 RMB per order, while the Logistics Cost Report (Essense Securities 2021) estimates that overall average labor cost (including base salaries, bonuses, etc.) of an order in last-mile delivery is 4.9 RMB per order at JD.com. Therefore, we estimate that the marginal labor cost of an order in last-mile delivery is between 2 RMB and 4.9 RMB. As for the electricity cost in last-mile delivery, the electricity consumption of delivering a package in last-mile delivery is around 0.0075 kWh (Carbonstop 2021). Given that the electricity price in China is 0.5 – 1 RMB per kWh, the electricity cost of a package in the last-mile delivery is estimated to be smaller than 0.01 RMB. Moreover, since this cost is typically borne by carriers, we assume it is absorbed into the labor cost.

Therefore, we estimate that the total variable transportation cost is between 2.4 ($= 0.4 + 2$) and 5.3 ($= 0.4 + 4.9$) RMB.

Packaging: The packaging cost refers to the expense of packaging materials (e.g., boxes, bags, padding materials) for packing an order, and thus it is considered as variable cost. The packaging cost may vary across orders depending on order sizes. According to Alibaba, the average cost of packaging per order lies between 1.5 RMB and 2 RMB. This range is in line with the estimated average packaging cost at JD.com, which is 1.7 RMB per order (Essense Securities 2021).

To sum up, we estimate that the marginal fulfillment cost of each online order is between 5.4 ($= 1.5 + 2.4 + 1.5$) RMB (0.8 USD) and 8.8 ($= 1.5 + 5.3 + 2.0$) RMB (1.3 USD). For simplicity, we take the average value 7.1 RMB (1.03 USD) as the estimate of the average marginal fulfillment cost of an additional online order.

B.2. Carbon Emissions

In general, according to the GHG (greenhouse gas) Protocol Corporate Standard, The GHG emissions of a company can be classified into three categories—Scope 1, 2 and 3 emissions (Greenhouse Gas Protocol 2004). Scope 1 emissions are those due to the company's direct operations, such as fuel combustion at company facilities or in company-owned cars. Scope 2 emissions are indirect emissions from the generation of purchased energy, and Scope 3 emissions are all indirect emissions (not included in scope 2) that occur in the value chain of the reporting company (US EPA 2020). Based on this framework, we can track the relevant GHG emissions each additional online order generates for the platform. Specifically, we consider the emissions from each of the three phases of an order: (i) order processing, (ii) packaging, and (iii) transportation. Note that,

consistent with the approach of Amazon (Amazon 2023), we do not take into account the carbon emissions from the manufacturing of the products that the platform sells nor from consumer use of these products.

Order Processing: The carbon emissions during the processing phase primarily stem from the electricity consumed by the platform’s office operations and by the platform’s data centers, which constitute Scope 2 emissions. The emissions from these infrastructures generally do not vary with order volume (unless capacity is exceeded). We have confirmed with Alibaba that the carbon emissions from the platform’s offices and data centers do not significantly increase with an order volume expansion at the scale of our earlier estimate (i.e., 15.00%). Thus, we consider the increase of emissions during the processing phase, due to the increase in orders, as negligible. Note that by making this assumption, our estimate constitutes a lower bound on the marginal carbon footprint impact of additional orders.

Packaging: Carbon emissions are released during the raw material extraction, processing, manufacturing, and disposal of packaging (Amazon 2023), constituting Scope 3 emissions for the platform. In the specific case of Tmall supermarket, orders are typically packed and delivered in corrugated boxes. Using a separate dataset we obtained from Alibaba, we estimate that the average package has a surface area of 0.36 square meters. We also learned from the platform that, based on their internal estimates, each gram of corrugated box is associated with 1.137g of GHG emissions. Hence, assuming an industry standard of 355 grams per square meter of corrugated board,¹⁴ we estimate that the packaging for each order results in $0.36 \times 355 \times 1.137 = 145.3$ g of GHG emissions on average. Note that this estimate does not consider changes in the average surface area of packaging with average basket size decreasing, as the estimated decrease in average basket size (1.14%) is too insignificant to necessitate a switch to smaller packaging boxes according to Alibaba’s current packaging operations. Specifically, available box sizes differ by at least 50% from each other. However, if the company offered wider box size selections such that the amount of packaging material used reduces at the same rate as the decrease in basket size, then the additional GHG emissions from each unit of packaging could instead be $145.3/(1 + 1.14\%) = 143.7$ g.

Transportation: Carbon emissions are released from transporting more orders, which can count as either Scope 1 or 3 emissions, depending on the mode and ownership of transportation. We again divide the transportation process of each order into two components: line-haul and last-mile delivery. Line-haul emissions are typically calculated by taking the product of shipment weight, distance traveled, and a relevant emission factor for the type of vehicle used (Blanco and Sheffi 2017). For our purposes, we use a carbon emission factor of 0.3764 kg CO₂e/(t·km) ($= \frac{0.5240+0.4211+0.1841}{3}$), which is the average of the carbon emission factors of three major types of trucks according to a recent report on carbon footprint in the Chinese e-commerce industry (Carbonstop 2021). We estimate the average travel distance of an order by using the average transportation distance by truck of all goods in China (i.e., 174 km), and use the average weight of all packages in China (i.e., 1.1336 kg) as the average weight of an order; see Carbonstop (2021) for a similar approach. Therefore, the

¹⁴ Tmall supermarkets use single-wall corrugated boxes, which consist of an outer liner, an inner liner and a corrugated flute. Following common industry standards, we assume the paper weights are 125 GSM (grams per square meter) for outer and inner liners and 105 GSM for the flute. Thus the total weight of the corrugated board is $125+125+105 = 355$ GSM. For more details on standard corrugated board configurations and weight calculations, please refer to <https://www.challengepackaging.co.uk/faqs/corrugated-board-grades-guide/>.

carbon emissions of a package during line-haul is estimated to be 74.2g ($=0.3764 \times 10^3 \times 1.1336 \times 10^{-3} \times 174$). Finally, for last-mile delivery emissions, we use the estimate in Carbonstop (2021), which assumes that an electric moped that delivers 200 packages a day charges 1.5 kWh every day. Thus the electricity consumption of delivering each package is 0.0075 kWh ($=\frac{1.5}{200}$), which translates into 4.5g of carbon emission per order on average.

Therefore, we estimate that the carbon emissions for each additional online order sum up to

$$4.5\text{g} + 74.2\text{g} + 145.3\text{g} = 224\text{g}.$$