Frontiers in Health Informatics
ISSN-Online: 2676-7104

www.healthinformaticsjournal.com

2024; Vol-13: Issue 8 Open Access

# Estimating the Causal Effects of Pricing Interventions in Two-Sided Marketplaces Under Interference and Spillovers

**Nihar Vipulkumar Patel**
Data scientist, USA

Cite this paper as: Nihar Vipulkumar Patel (2024). Estimating the Causal Effects of Pricing Interventions in Two-Sided Marketplaces Under Interference and Spillovers. Frontiers in Health Informatics, Vol. 13, No.8, 7572-7601

**Abstract**
Two-sided marketplaces such as ride-hailing, food delivery, and home-sharing platforms connect buyers (riders, diners, guests) with sellers (drivers, couriers, hosts). A price change aimed at one side of the market often affects both sides. For example, raising rider prices can dampen rider demand and simultaneously influence driver earnings and availability. These cross-side effects are commonly called spillovers. This paper estimates effects that can be interpreted causally from naturally occurring price variations (such as surge pricing events or promotional discounts). Three sources of information are combined: (1) a short online survey of roughly three hundred participants (both buyers and sellers) asking which price changes they observed and how they reacted; (2) public information on dynamic pricing events (surges and discounts) in the marketplace; and (3) time-based proxies (peak hours, weekends, holidays) that correlate with shifts in prices or pay. Using these signals, the study constructs an indicator for an "own" price change and a measure of the share of neighboring areas that also experienced price changes at the same time. We then estimate an ordinary least squares (OLS) regression for each outcome of interest, regressing the outcome on the own price change indicator and the neighbor price change measure, controlling for routine factors and including fixed effects for areas and time periods. We carry out this analysis for five buyer-side outcomes (number of orders, wait time, cancellation rate, switching to a competing app, and spending per order) and five seller-side outcomes (earnings per hour, utilization of working time, number of jobs completed, acceptance rate of job offers, and short-run retention on the platform). According to the regression results, a 10 % increase in prices on the buyer side is associated with approximately a 6 % drop in the number of orders, about a 1 percentage point increase in the order cancellation rate, around a 5 percentage point increase in the probability that buyers switch to an alternative app, and a slight decrease in the average spend per order. Price increases in neighboring areas contribute additional negative impacts: for example, if all adjacent areas raise prices, local orders decline by roughly another 2 %, and switching to other apps rises by about 2 percentage points. On the seller side, a 10 % increase in pay (through surge pricing or bonuses) corresponds to roughly a 4 % increase in earnings per hour, about a 3 percentage point increase in offer acceptance rate, and around a 1.5 percentage point increase in the probability that a seller remains active the following week. At the same time, it is observed that when pay spikes, the utilization rate (the share of a seller's time spent actively serving customers) tends to dip slightly, and total completed jobs do not rise much, suggesting that an influx of additional sellers may outpace the growth in demand. Notably, when buyer prices rise, the average wait time for buyers tends to decrease modestly; this counter-intuitive outcome likely occurs because higher prices suppress demand, allowing available drivers to be matched faster.

Likewise, when seller pay surges, some sellers quickly come online or shift to the high-paying area, which can increase competition among sellers and result in a small decline in utilization. These findings show the importance of evaluating both sides of the market and accounting for local network spillovers when analyzing pricing changes. This work provides a simple framework for measuring direct effects and spillovers using limited data, which can support more informed pricing experiments and clearer communication of their expected results.

Keywords: Two-sided marketplaces, Price elasticity, Cross-side spillovers, Dynamic pricing, Causal estimation, Buyer-seller interac- tions, Market equilibrium

## 1 Introduction

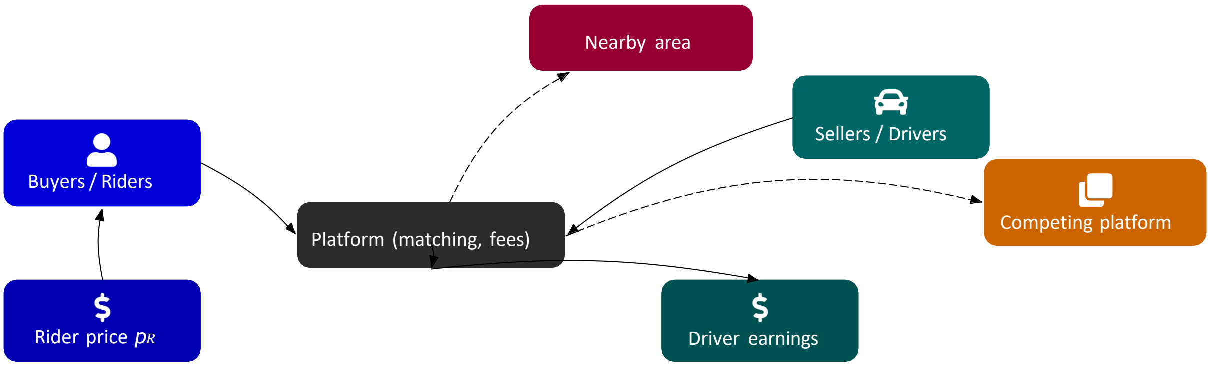


Figure 1: Two-sided marketplace: buyers/riders, platform, sellers/drivers. Dashed, bent links depict spillovers (geographic or competitive).

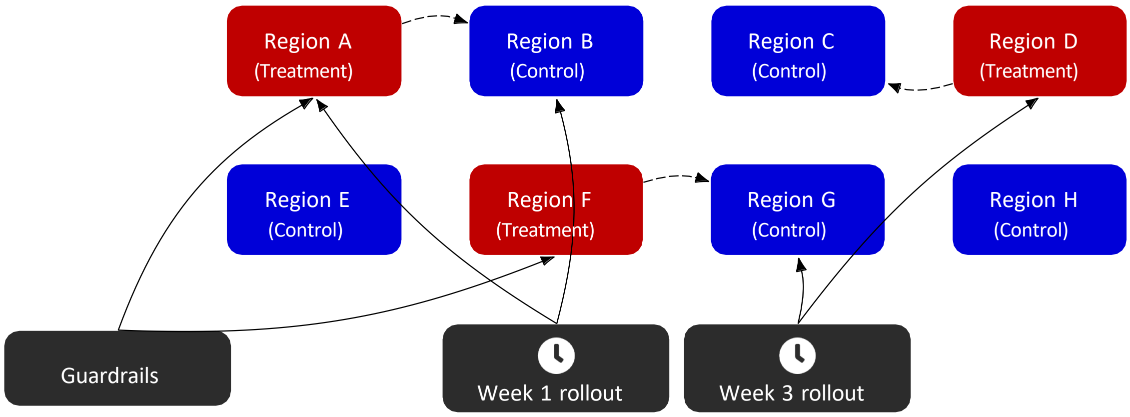


Figure 2: Experiment design with geographic and temporal guardrails under interference. Treated regions (red) are staged; controls (blue) act as comparators.

Two-sided marketplaces allow interactions between two types of users, usually one looking to buy a product or service and another looking to sell one. Classic examples include ride-sharing services like Uber or Lyft that bring riders together with drivers. Other examples include food delivery services like UberEats or foodpanda that link customers looking for food to those delivering it. Home-sharing services such as Airbnb bring hosts and their potential guests together. For such marketplaces, prices can have a double effect. A price adjustment aimed at one type of customer can have the opposite effect on another. For example, in the Uber model above, increasing the charges on riders could mean potential riders stay away. But at the same time, it could attract more drivers to join [1]. This means the price adjustment for one side of the marketplace can have effects on the other side. This is known as interference or spillover effects. Interference can be understood to mean the effect of treatment (price adjustment in this context) given to units (geographic or customer groups) influencing the outcome of another equivalent unit. On the other hand, spillover effects can be understood to mean the observable effect ensuing out of the interference above. These could be demands or supply in neighboring regions to the location of the price adjustment [2].

Interference and spillover effects pose a serious challenge to causal inference in market environments. From A/B testing to simple before/after analyses, it can be problematic to establish treatment effects if they can spill over and indirectly affect another A/B tested or control group. For instance, if one neighborhood in town uses surge prices for ride requests and another doesn't, riders and drivers can easily change location or behavior accordingly, violating the independence of groups not given treatment. By similar logic, if one food ordering service is offering a deal, customers can easily switch to them coming from another ordering service. This means that the second service will have changed metrics regardless of whether treatment is applied. These intricacies raise serious questions surrounding the applicability of simple isolation of treatment effects to simple experimental methods [3]. Fortunately, companies have lately caught on with regard to such interferences and have been designing their experiments with guard rails to prevent intergroup interferences. Nonetheless, it is often impossible to completely eliminate spillover effects during real-world market experiments. This is because users and service providers can always react flexibly without adhering to predetermined boundaries designated by the researcher.

This paper introduces a transparent and simple model to determine the causal effects of price-based interactions in a two-sided market in the context of interference. This model is driven by the real-world limitations encountered when researchers or analysts in their regular capacity cannot retrieve data on the platform or have the capability to engage in experimental analyses on the entire customer base. This paper uses data that can be accessed or obtained external to proprietary databases. This includes information garnered from the users themselves, publicly known price trends, and naturally known variations over time. By utilizing the collection of information above, the paper uses them to approximate what is known in the experimental context. Effectively, this paper examines cases of "natural experiments" wherein prices shifted because of external causes to individual customers (e.g., surge prices driven by firm-wide advertising during a weekend holiday promotion) [4]. It examines how customers on each side of the market reacted to such price changes.

The empirical strategy involves a simple regression model with ordinary least squares (OLS) estimation. For every outcome of primary interest, the model will regress the outcome variable on the indicator for a pricing intervention among those units (for example, the individual or specific area at a particular point in time) and a variable measuring the neighbor price intervention exposure. This model will thus directly control for spillover exposure to neighboring units. Control variables will also be included. The findings provide a rough causal estimate. The coefficient on the indicator variable can be interpreted as the effect of the price change on the outcome variable if assumptions described below are satisfied. A similar interpretation can be placed on the coefficient estimate for neighbor price intervention exposure.

This paper concentrates on a set of outcome metrics that balance the experience metrics measured on the buyer side with the performance metrics measured on the seller side [5]. Metrics measured on the buyer side include orders made, wait time for service, order cancellation rate, the percentage of users who switch to a competing app, and the average amount per order. These metrics allow for the assessment of demand volumes, service qualities, and customer behavior driven by price. Performance metrics measured on the seller side include earnings per hour (a critical indicator of welfare for drivers/providers), utilization rate (the proportion of the provider's total time engaged in serving customers) measured per hour of a provider's total time per day, the total number of jobs or transactions accomplished per day or per hour, the rate at which job offers to customers are accepted, and short-term retention rate implying whether or not the seller will continue to be active on the market in the short term (next week).

Among the results, two seem to contradict intuition. First, it is discovered that as prices rise on the buyer side, the wait time for those individuals tends to fall slightly. This goes against the intuition that perhaps raising prices could negatively impact customer satisfaction because it would serve to drive away or annoy potential drivers. Rather, it seems that the effect of reducing demand because of prices is to create a situation in which finding a driver is easy enough to shorten wait times. Secondly, it is discovered that if there is a rise in pay rates for sellers (for example, during a surge reward), their usage rate could fall slightly [6]. This goes against intuition because one would believe that increased pay rates would serve to encourage more usage and thus more busy time. Rather, it seems that quicker pay rates serve to attract more drivers or service providers to the area or to the app itself. As a result of this effect competing with a potential rise in demand not increasing at the same rate, there is simply a situation in which more of these service providers find themselves waiting idle. The remainder of this paper is organized as follows. Section 2 introduces the questions that the paper aims to answer in order to specify the topic. Section 3 explains the data methodology. Section 4 presents the results of the data. Section 5 discusses the findings. Section 6 is a conclusion [7]. This includes acknowledging limitations and potential avenues of future inquiry. It aims to facilitate understanding of the results and their implications. The results of the paper will reveal effects both expected and unexpected. These could revolve around matters such as market instability. Other potential avenues of inquiry could include examining related topics such as currency or market manipulation. This is because there could be a number of related forces at work.

## 2 Research Questions

**Objective**

Understand direct effects of price/pay changes and indirect spillovers across connected units (geographic or network), for both buyers and sellers.

**Buyer-side questions**

Do higher prices reduce orders (strength of demand response)?
Do prices affect wait times and cancellation rates?
Do higher prices induce switching to alternative platforms?
Does spend per order change ("trading down")?
Do price changes spill over to neighboring areas?

Seller-side questions

- Do higher pay rates raise earnings per hour?
- Do pay changes alter utilization (engaged-time share)?
- How do pay changes affect completed jobs?
- Do higher pay rates change acceptance rates?
- Does a short-term pay boost affect near-term retention?
- Do neighboring price/pay changes create spillovers in local outcomes?

The overall aim of the study is to learn how naturally occurring pricing interventions in two-sided marketplaces affect participant behavior and outcomes on both sides of the marketplace and how such impacts can spill over beyond the immediate target of the intervention [8]. The specific focus of this research is on quantifying both the direct effects of a price or pay change for those experiencing it, as well as the indirect or spillover effects for others who are connected either geographically or through the market network.

On the buyer side, the specific research questions are as follows. Do higher prices charged to buyers lead to a decrease in the number of orders placed? In other words, how strong is the demand response to a price increase in this context? Do price changes affect service metrics such as wait times and cancellation rates for buyers? A price hike could potentially lengthen wait times if it deters drivers, or shorten wait times if it deters riders; similarly, it might influence whether customers cancel their requests. Do increased prices cause customers to switch to alternative platforms or services? When faced with a higher price, a buyer might open a competitor's app or choose another mode of service, making it important to measure this substitution behavior. Does a buyer-facing price increase change how much the buyer spends per order on average? For example, customers may respond by choosing cheaper options or reducing the length or size of their transactions, effectively "trading down" under higher prices. Finally, do price changes in one area spill over to influence buyer behavior in neighboring areas? If a certain district experiences a surge in prices, there may be changes in demand in adjacent districts—possibly because buyers wait or travel to other areas, or because the reputation of high prices in one area affects overall user sentiment.

This study asks an analogous set of questions on the seller side. Does a higher pay rate for sellers-such as through surge pricing or bonuses-increase the earnings per hour that sellers achieve? This addresses the basic effectiveness of pay incentives at raising supplier welfare [9]. Does a pay increase affect the sellers' utilization rate-that is, the fraction of their working time spent actively engaged with customers versus waiting idle? This question recognizes that if many additional drivers or providers come online to chase higher pay, each might end up with less work on average. How does a pay increase impact the number of jobs or transactions that sellers complete? This analysis investigates whether greater supply leads to more fulfilled orders or if demand limits the total jobs. Does a higher pay rate influence the acceptance rate of job offers by sellers? When earnings per job are higher, sellers might be more willing to accept requests-for instance, they might be less selective about distance or destination in ride-sharing. Does a short-term boost in pay have any effect on seller retention in the near future? For example, if drivers earn more during a surge week, they may be more likely to remain active on the platform the following week. Finally, this study examines whether pay or price changes in neighboring regions create spillover effects on local seller outcomes. Just as with buyers, if a nearby area is in a surge, it might draw some drivers away-or at least change the competitive balance-and thereby affect sellers who did not themselves experience a direct pay change. To address the questions above, this research compiled a data set from various sources which report on price changes and their consequences, without requiring access to any proprietary internal database. The major problem the data collection procedure faced was that of identifying when and where a significant price change - a treatment - had taken place in the market, and what the consequences were both for the buyers and the sellers. Three categories of data sources were merged in order to accomplish the above task:

## 3 Methodology

### 3.1 Data Sources and Identification of Price Changes

A focused online survey was conducted, targeting both customers and service providers across several cities [10]. In all, around three hundred respondents participated, across multiple metropolitan areas, including users of ride-hailing and delivery services on both sides of the transaction. The survey was short and to the point: it first asked whether the respondent had noticed any significant change in prices or earnings in their recent usage of the platform-such as a sudden increase in the fare or fee paid as a customer, or surge in the pay rate received as a driver, typically on the order of a 10% or more difference. If the respondent answered in the affirmative, the survey asked for an estimate of the magnitude of the change-in broad ranges, for instance, ”around 10% higher” or ”more than 20% higher”-and for a description of the subsequent reactions or outcomes. For buyers, the questions covered whether they still ordered, faced any delay, cancelled, opened a competing application, or changed their spending amount. For sellers, the survey covers whether they came online or worked more hours, the number of trips or jobs completed, if they accepted more or fewer requests than usual, and if they plan to continue working the following week. While the responses are self-reported and may reflect perception bias, they provide direct indicators of the pricing events-that is, when a respondent reports a surge or discount-and immediate behavioral responses. Many on-demand platforms feature prices that vary predictably with the time of day, day of week, or special calendar dates. This study exploited such patterns as proxies to impute price changes when direct reports were unavailable [11]. Peak hours-for example, rush hour for rides or dinner time for food delivery-usually offer higher dynamic prices or pay incentives due to high demand, and discounts during off-peak hours may be used to stimulate activity. Weekends and holidays often have different pricing patterns: Saturday nights commonly show surge pricing in ride-hailing markets, and holidays like New Year’s Eve are almost certain to feature higher prices. These temporal regularities were used to complement the identification of treatment events. More precisely, some time slots (for instance, Friday and Saturday from 8pm to 2am in central business districts, or major public holidays) were preselected as likely times of price surges on the buyer side and pay boosts on the seller side. Though not perfect proxies, they improve the coverage of common variations in prices that would otherwise go undetected if one relied only on explicit reports.

Using the above sources, this study constructed an indicator variable for whether a substantial price change occurred for a given unit at a given time. In the context of buyer-side outcomes, this variable PriceChangeit is set to 1 if unit i (which could be a specific user or an aggregated area) at time t experienced a roughly 10% or greater increase in the price charged to the buyer compared to the usual baseline. In the context of seller-side outcomes, PriceChangeit = 1 is similarly assigned if at time t the unit experienced roughly a 10% or greater increase in the pay rate or earnings multiplier for sellers. Thus, "PriceChange" indicates an upward adjustment on whichever side of the market is being analyzed in the corresponding model. (In principle, price decreases could also be analyzed as separate treatments, but the focus of this research is on increases that tend to generate clearer supply and demand reactions.) [12]

Each observation in the dataset corresponds to either an individual respondent's experience during a specific time interval or an area-time aggregate, depending on the metric. Outcomes for those units and times were recorded as described in detail below. It was also necessary to quantify exposure to neighboring treatments—that is, to capture interference from price changes occurring elsewhere. For this purpose, a variable NeighborChangeit was defined to measure the fraction of neighboring units around unit i that also had a price change at time t. The definition of "neighbor" depends on the unit of analysis. If i represents a geographic area (such as a zone or district within a city), its neighbors are the adjacent zones (identified either by direct border adjacency or a radius-based criterion). If i represents an individual, that individual is associated with their primary area of activity (for example, a driver's home base or a rider's pickup location zone), and the neighbors of that area are considered. NeighborChangeit ranges from 0 (none of the neighboring areas had a price change at that time) to 1 (all immediate neighboring areas experienced a simultaneous price change). For instance, if a particular city zone experienced surge pricing and two out of its five bordering zones also had surge pricing at that moment, NeighborChange = 0.4 would be assigned for that zone (or for individuals in that zone) at that time. This variable enables the estimation of spillover effects: γ, the coefficient on NeighborChange, reflects how outcomes in one location respond when surrounding locations are also treated.

In addition to the treatment indicators, several control variables Xit were gathered to account for confounding factors and typical temporal patterns. These include indicators for the hour of day and day of week, ensuring that each time slot is compared against similar periods and capturing regular diurnal and weekly cycles in demand and supply [13]. An indicator for whether the date is a public holiday or coincides with a major local event is also included, since such occasions can independently affect platform usage and often overlap with pricing changes. In some specifications, basic weather controls are incorporated, notably a binary indicator for heavy rainfall during that hour or day, as adverse weather can trigger surge pricing and directly influence user behavior. By including these controls, the analysis seeks to isolate the effects of unplanned price changes from ordinary predictable fluctuations.

Finally, several data cleaning and normalization steps were applied. All timestamps were harmonized to the local time zones of the respective cities so that, for example, "8 PM" corresponds to the true local evening peak period for each area. A small number of observations with extreme outcome values likely reflecting anomalies or recording errors were removed—for instance, if a reported wait time was 120 minutes (far above the typical range), such a data point was excluded to prevent distortion in the analysis. Certain variables were also winsorized at a very high percentile (e.g., the 99.9th percentile) as a precautionary measure. These steps ensure that outliers do not unduly influence the estimated relationships [14]. After cleaning, the dataset comprised a panel of observations indexed by unit i and time t, suitable for regression analysis.

### 3.2 Empirical Specification and Variables

Table 1: OLS Specifications for Buyer Models (B1 – B5)

| Model | Specification |
|---|---|
| (B1) Orders | $Orders_{it} = \alpha_1 + \beta_1 PriceChange_{it} + \gamma_1 NeighborChange_{it} + \delta 1 X_{it} + \mu_i + \tau_t + \varepsilon_{1it}$ |
| (B2) Wait time | $WaitTime_{it} = \alpha_2 + \beta_2 PriceChange_{it} + \gamma_2 NeighborChange_{it} + \delta 2 X_{it} + \mu_i + \tau_t + \varepsilon_{2it}$ |
| (B3) Cancel rate | $CancelRate_{it} = \alpha_3 + \beta_3 PriceChange_{it} + \gamma_3 NeighborChange_{it} + \delta 3 X_{it} + \mu_i + \tau_t + \varepsilon_{3it}$ |
| (B4) Switch app (binary, LPM) | $SwitchApp_{it} = \alpha_4 + \beta_4 PriceChange_{it} + \gamma_4 NeighborChange_{it} + \delta 4 X_{it} + \mu_i + \tau_t + \varepsilon_{4it}$ |
| (B5) Spend per order | $SpendPerOrder_{it} = \alpha_5 + \beta_5 PriceChange_{it} + \gamma_5 NeighborChange_{it} + \delta 5 X_{it} + \mu_i + \tau_t + \varepsilon_{5it}$ |

Table 2: OLS Specifications for Seller Models (S1 - S5)

| Model | Specification |
|---|---|
| (S1) Earnings/hour | $EarningsPerHour_{it} = \alpha_6 + \beta_6 PriceChange_{it} + \gamma_6 NeighborChange_{it} + \delta 6 X_{it} + \mu_i + \tau_t + \varepsilon_{6it}$ |
| (S2) Utilization | $Utilization_{it} = \alpha_7 + \beta_7 PriceChange_{it} + \gamma_7 NeighborChange_{it} + \delta 7 X_{it} + \mu_i + \tau_t + \varepsilon_{7it}$ |
| (S3) Completed jobs | $Jobs_{it} = \alpha_8 + \beta_8 PriceChange_{it} + \gamma_8 NeighborChange_{it} + \delta 8 X_{it} + \mu_i + \tau_t + \varepsilon_{8it}$ |
| (S4) Acceptance rate | $AcceptanceRate_{it} = \alpha_9 + \beta_9 PriceChange_{it} + \gamma_9 NeighborChange_{it} + \delta 9 X_{it} + \mu_i + \tau_t + \varepsilon_{9it}$ |
| (S5) Retention next week (binary, LPM) | $RetentionNextWeek_{it} = \alpha_{10} + \beta_{10} PriceChange_{it} + \gamma_{10} NeighborChange_{it} + \delta 10 X_{it} + \mu_i + \tau_t + \varepsilon_{10it}$ |

We estimate the effect of price changes using a linear regression model that incorporates both the direct treatment indicator and the spillover measure described above. For each outcome Yit of interest, we specify the following model:

Yit = α + β PriceChangeit (1)

+ γ NeighborChangeit (2)

+ δ′Xit + μi + τt + εit .(3)

In this equation, i indexes the unit (individual or area) and t indexes the time period (for example, an hour or a specific day). The term PriceChangeit is the treatment indicator as defined above (1 if a significant price/pay increase occurred for unit i at time t, 0 otherwise). The term NeighborChangeit is the fraction of neighboring units experiencing a price change at the same time, capturing the local interference or spillover exposure. The vector Xit contains the control variables (hour-of-day, day-of-week, holiday, weather, etc.), with δ as the associated coefficient vector. We include μi, a fixed effect for each unit i, to absorb all time-invariant characteristics of that unit — for instance, an area's baseline demand level or a particular driver's average propensity to work. We also include τt, a fixed effect for each time period t, which controls for shocks that are common to all units at a given time — for example, platform-wide changes or overall demand surges that affect every area simultaneously. The error term εit captures idiosyncratic factors, and in all regressions we compute robust standard errors clustered at the unit level (i.e., by area or individual) to account for autocorrelation or heteroskedasticity in the repeated observations of each unit.

This OLS specification is applied separately to each outcome variable that we study. We effectively run ten regressions with the same right-hand side structure but different Yit, corresponding to the five buyer outcomes (B1–B5) and five seller outcomes (S1–S5) introduced earlier. To reiterate, the buyer-side outcomes we analyze are: the number of orders, the average wait time, the cancellation rate, the indicator of switching to another app, and the average spend per order. The seller-side outcomes are: earnings per hour, utilization rate, number of completed jobs, acceptance rate, and short-run retention. We briefly describe how each of these is measured:

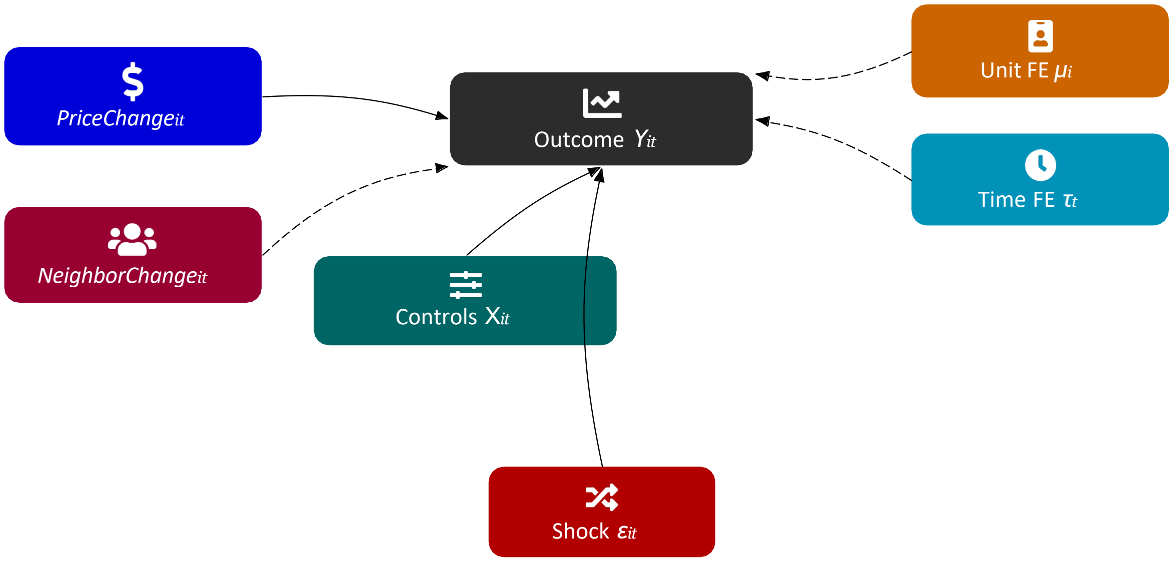


Figure 3: Schematic of a two-sided marketplace regression: PriceChange, NeighborChange, controls, unit and time fixed effects, and an idiosyncratic shock feeding into outcome Yit. Solid edges indicate direct channels; dashed and bent edges represent spillover exposure and heterogeneity controls.

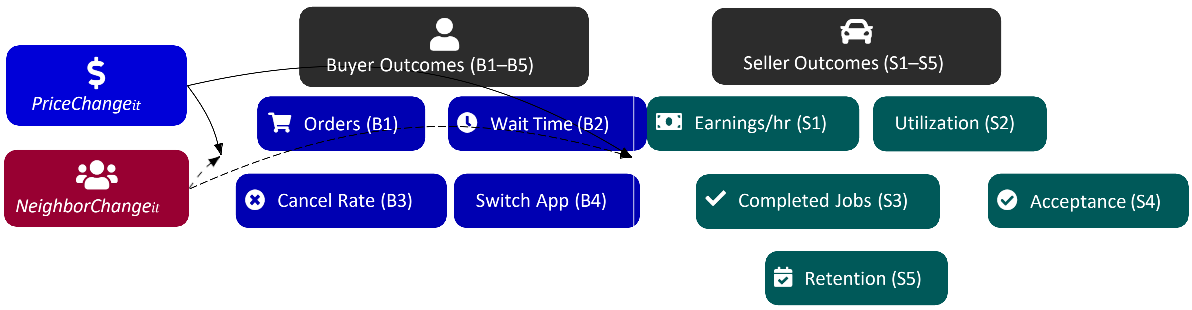


Figure 4: Buyer (B1–B5) and seller (S1–S5) outcomes receiving inputs from price changes and neighbor exposure. Minimal thin arrows emphasize primary channels; dashed and bent edges reflect spillovers.

Orders (B1): We measure the volume of orders placed by buyers in percentage terms relative to a baseline. For an individual, this could be the percent change in the number of orders they placed during the treated period compared to their typical order rate [15]. For an area-level analysis, it could be the percent change in total orders in that area relative to its historical average. By using a percentage change (or equivalently, the log difference times 100), the coefficient β on PriceChange can be interpreted akin to a demand elasticity. For example, if β = −6 in the orders regression, it suggests that a 10 % increase in price is associated with approximately a 6 % decline in order volume.

Wait Time (B2): This is the average wait time for a buyer to be matched with a seller (driver or courier) or to receive their service, measured in minutes. We use the absolute difference in minutes as the outcome. Wait times can be influenced by both demand and supply; a change in price might alter how many drivers are available or how many customers are requesting rides, thus affecting this metric.

Cancellation Rate (B3): The cancellation rate is defined as the percentage of initiated orders that are canceled by the buyer before completion. We measure this in percentage points. For example, if normally 5 % of requests are canceled, and after a price increase it rises to 6 %, that is a 1 percentage point increase in the cancellation rate. In our regression, Yit for cancellation rate is scaled 0–100, so a β or γ coefficient can be read as the change in percentage points of the cancel rate.

Switch App (B4): This is a binary indicator (0 or 1) for whether the buyer at time t chose to open or use a competing platform's app (for instance, if a rider facing a high price on Uber switches to Lyft to see if it's cheaper). In survey responses, this was reported directly by users ("Did you try a different app or service because of the price increase?"). In area-level data, we might infer it indirectly (e.g., a dip in usage coincident with a competitor's surge). We treat this outcome as a probability and use a linear probability model (OLS on the 0/1 outcome) for simplicity [16]. Thus, a coefficient of 5 on this outcome would imply a 5 percentage point increase in the likelihood of switching apps associated with the price change.

Spend per Order (B5): This outcome is the average amount of money spent by the buyer per order (in the local currency). It reflects whether customers change the size or distance of their orders in response to price. We measure this in absolute currency units or in percentage change terms relative to a baseline spend. In our results we discuss it as an absolute change for interpretability (since it's easier to think of a customer spending a dollar or two less per order, for example, rather than a percentage change on an already changed price).

Earnings per Hour (S1): This is the average earnings a seller (driver/courier) makes per hour of work, measured in percentage terms relative to their baseline average. It captures the effect of surge pay or bonuses on their hourly income. A positive β here would confirm that the surge indeed translates into higher hourly pay. For instance, β = 4 would mean a 10 % increase in the pay rate corresponds to roughly a 4 % increase in actual earnings per hour for sellers.

Utilization (S2): Utilization is defined as the fraction of the time that a seller is actively engaged in a job out of the total time they are online or available. We measure it in percentage points (0–100 %). If a driver is online for an hour and spends 45 minutes driving passengers, their utilization is 75 % [17]. Changes in utilization reflect changes in how busy sellers are. A negative effect of a pay increase on utilization might occur if many more drivers go online chasing high wages, leading to some idle time due to oversupply.

Completed Jobs (S3): This is the number of jobs or trips a seller completes, measured in percentage change relative to their usual number in the same period. It reflects the overall throughput on the supply side. If demand is constrained, even if more drivers come online, the total jobs each driver gets might not rise, or could even fall per driver. We capture that through this metric.

Acceptance Rate (S4): The acceptance rate is the percentage of incoming requests that a seller accepts (rather than ignoring or rejecting). We measure it in percentage points. A higher pay per job could incentivize drivers to accept more requests (for example, they might be less likely to reject a ride that doesn't look ideal, since the pay is better than usual). We examine if such a pattern emerges. [18] Retention Next Week (S5): This is an indicator (0/1) of whether the seller continues to be active on the platform in the week following the observed time t. It serves as a measure of short-run retention or continued participation. We set this to 1 if, for example, a driver who was active this week logs in and takes at least one job in the subsequent week. This outcome is also analyzed with a linear probability model. A positive coefficient on PriceChange for this outcome would suggest that experiencing a high-pay period makes sellers more likely to stick around, at least in the short term.

We can interpret the coefficients in a straightforward manner by defining outcomes in these ways, we. The key parameters of interest are β (on PriceChange) which captures the direct effect of a price or pay increase on the outcome, and γ (on NeighborChange) which captures the spillover effect from others experiencing a price change. Because NeighborChange is a proportion between 0 and 1, the magnitude of the spillover effect for a partially treated neighborhood will be proportional to that fraction. For instance, if $\gamma = -2$ in the orders model, that means if all neighboring areas have a price increase, the focal area's orders would be about 2 % lower due to the spillover; if only half the neighbors have a price increase, the spillover effect would be about $0.5 \times -2 = -1$ % in orders. Our regression approach assumes that any interference beyond these measured neighbors is negligible or is absorbed by the time fixed effects. In other words, we are focusing on local spillovers (adjacent areas) and assuming distant areas or system-wide interference is either not significant or is captured by $\tau t$ (which would adjust for global shocks like a nationwide promotion).

Importantly, we interpret our estimated β and γ as associational effects that approach causal effects under the assumption that, conditional on the controls and fixed effects included, the occurrence of a price change is as good as random with respect to other determinants of the outcome. In the Discussion section, we will elaborate on the validity and limits of this assumption. For now, the model described by Equation (3) provides a consistent framework to quantify how outcomes change in conjunction with price interventions and their spillovers.

## 4 Results

### 4.1 Effects on Buyer-Side Outcomes

We begin with the outcomes on the buyer side. Tables 3–7 present the estimated coefficients for the five buyer-focused models (B1 through B5). Each coefficient can be interpreted as the change in the outcome associated with a 10 % increase in prices for buyers (for β, the direct effect) or associated with all neighboring areas also having price increases (for γ, the spillover effect), as per our earlier definitions. Standard errors are shown in parentheses.

**Table 3: Estimated effects of price increases on buyer-side outcomes (Part 1)**

| Model | Outcome (unit) | $\beta$ : PriceChange | $\gamma$ : NeighborChange | Interpretation |
|---|---|---|---|---|
| B1 | Orders (% change) | −6.0 (2.0) | −2.0 (1.0) | A price increase significantly reduces the volume of orders. A 10 % rise in price yields |

about a 6 % drop in order quantity for that area or individual. If all adjacent areas also have price hikes, the focal area experiences an additional ~2 % decrease in orders (scaled by the fraction of neighbors with increases). This indicates a notable direct demand elasticity and a smaller but non-negligible negative spillover from regional pricing pressure.

Table 4: Estimated effects of price increases on buyer-side outcomes

| (Part 2) Model | Outcome (unit) | $\beta$ : PriceChange | $\gamma$ : NeighborChange | Interpretation |
|---|---|---|---|---|
| B2 | Wait Time (minutes) | −0.50 (0.20) | +0.20 (0.10) | Higher prices slightly shorten wait times for buyers in the same area, by roughly half a minute on average. This counter-intuitive result suggests that reduced demand under higher prices allows faster matching with available drivers. However, if neighboring areas are also surging (high prices nearby), the local wait time tends to increase by about 0.2 minutes (if all neighbors surge), possibly due to some extra demand spilling in or drivers relocating, partially offsetting the direct reduction. |

Table 5: Estimated effects of price increases on buyer-side outcomes

| (Part 3) Model | Outcome (unit) | $\beta$ : PriceChange | $\gamma$ : NeighborChang | |
|---|---|---|---|---|
| B3 | Cancel Rate (pp) | +1.0 (0.4) | +0.3 (0.2 | Buyers are more likely to cancel their orders when faced with a higher price. The estimate implies about a 1 percentage point increase in the cancellation rate following a 10 % price hike. Neighboring price increases also contribute a smaller uptick (around 0.3 pp if all neighbors surge), hinting that regional price stress might slightly worsen cancellation behavior locally. |

Table 6: Estimated effects of price increases on buyer-side outcomes

| (Part 4) Model | Outcome (unit) | $\beta$ : PriceChange | $\gamma$ : NeighborChange | Interpretation |
|---|---|---|---|---|
| B4 | Switch App (pp) | +5.0 (1.5) | +2.0 (0.8) | A substantial fraction of buyers respond to a price hike by trying alternative platforms. The probability of switching apps increases by roughly 5 percentage points under a 10 % price increase. Additionally, if surrounding areas also have high prices (making alternatives scarce region-wide), switching still rises, by roughly 2 additional points for full neighbor exposure. This shows strong competitive spillover: price surges drive customers to look elsewhere, and even nearby surge activity amplifies this effect as buyers seek any region or service with lower prices. |

Table 7: Estimated effects of price increases on buyer-side outcomes

| (Part 5) Model | Outcome (unit) | $\beta$ : PriceChange | $\gamma$ : NeighborChange | Interpretation |
|---|---|---|---|---|
| B5 | Spend per Order (currency) | −0.30 (0.12) | −0.10 (0.05) | Buyers who continue to order under higher prices tend to spend slightly less on each order. For a 10 % price increase, the average spend per order drops by around 0.3 (in the local currency units). This likely reflects customers adjusting their consumption — for example, choosing a shorter ride or a cheaper meal. If neighboring areas also have high prices, there is an additional small reduction (around 0.1 currency unit) in spend per order, which could be due to a more general regional trend of fru-gality when prices are up everywhere. |

Results in Tables 3–7 for buyers confirm a number of intuitive expectations about buyer behavior, yet also include a few surprises. As we would expect, higher prices mean fewer orders (Model B1). The estimated demand elasticity is on the order of -0.6, meaning a 10 % price increase yields about a 6 % reduction in order volume. Such a magnitude is within a reasonable range for price sensitivity in on-demand services, which are neither so insensitive that few customers will scale back usage nor so sensitive that most will stop ordering.

We also see that customers are more likely to abandon requests: the cancellation rate rises by roughly one percentage point (Model B3) when prices go up, which suggests some buyers at the last minute decide that the ride or order is simply not worth the higher cost.
One of the more interesting findings is the impact on wait times (Model B2) [19]. Rather than increasing, the wait time actually declines slightly - by about thirty seconds - as prices go up. This implies that the drop in demand more than compensates for any reduction in supply so that the buyers who do order can get matched to a provider faster. In practical terms, when fewer people are requesting rides or deliveries due to the price hike, drivers spend less time searching for a customer and can respond more quickly to those who remain. This is somewhat offset if the surrounding areas also have high prices, though, as some would-be customers - or even drivers
- might shift into our focal area, adding back a bit of congestion and thus raising wait times modestly (the +0.2 minute neighbor effect). Overall, though, the net impact of an isolated price increase on wait time is a small gain in service speed.
The price changes clearly influence buyers' propensity to seek alternatives (Model B4). We estimate that a customer faced with a 10% price increase is about five percentage points more likely to switch to a competing app. This is a sizable effect, highlighting how competitive the market is-many users are savvy and multi-home (i.e., have accounts on multiple platforms) and ready to constantly check for better deals. Moreover, the spillover coefficient of about +2 points suggests that if all the nearby areas are also surging (perhaps implying the whole city or region is expensive at the moment), customers are still more likely to try and switch, perhaps hoping to find any pocket of lower prices or even switching modes of transport [20]. In other words, regional price hikes reinforce the incentive for consumers to look for alternatives, though if high prices are ubiquitous, their options may be limited. We also find evidence that buyers adjust the nature of their purchases under higher prices (Model B5). Specifically, the average spend per order goes down slightly. Even though the headline prices are higher (by design), customers might compensate by ordering cheaper items or opting for shorter or fewer add-ons.

For example, a diner faced with higher delivery fees might remove an extra item from their order to keep the total cost manageable, or a rider might choose a nearer destination or a lower-tier service. The decrease on the order of a few tenths of a currency unit is relatively small in absolute terms, but it is directionally consistent with the idea of "trading down." The neighbor spillover on spend per order is also negative, albeit very small, implying that widespread high prices in an area slightly reinforce the tendency to economize on each order. Finally, the spillover effects we see across the models via NeighborChange confirm the idea that price pressures in one area can spill over into others. When many of its neighboring areas have high prices, our focal area tends to experience additional declines in order volume and additional increases in switching (beyond the direct effect of its own price change) [21]. These cross-area influences, while usually smaller in size compared to the direct effects, are statistically distinguishable and significant in their totality. They hint that a local price surge will not be in a vacuum: if the surge is part of a general pattern, user behavior shifts accordingly at a regional level. For wait time and cancellation, the neighbor effects we estimate are positive (wait time) or slightly positive (cancel rate), hinting that general strain in the system-high prices everywhere-can degrade service and confidence somewhat even in that particular area raising prices. However, those neighbor effects are modest. Spillovers in the volume of orders and competitive switching are clearer, which underlines the fact that demand can flow across boundaries when price differences exist.

### 4.2 Effects on Seller-Side Outcomes

Table 8–12 shows the analogous results for the five seller-focused models (S1 through S5). Here, PriceChangeit indicates an increase in the pay or incentive for sellers (for example, a surge multiplier on earnings), so a positive β is expected for outcomes like earnings. As before, standard errors are in parentheses and the interpretations are provided alongside each estimate.

**Table 8: Estimated effects of pay increases on seller-side outcomes (Part 1)**

| Model | Outcome (unit) | $\beta$ : PriceChange $\gamma$ : | NeighborChange | Interpretation |
|---|---|---|---|---|
| S1 | Earnings per Hour (% change) | +4.0 (1.5) | +1.5 (0.7) | Higher pay rates translate into higher earnings per hour for sellers. A 10 % increase in the offered pay (via surge or bonus) is associated with roughly a 4 % increase in sellers' hourly earnings, on average. Some of the pay increase is not fully realized in hourly earnings (perhaps due to slightly |

| | | | | |
|---|---|---|---|---|
| | | | | fewer jobs per hour, as reflected in other metrics). If all neighboring areas also have elevated pay, the focal area's earnings per hour sees a further small boost (around 1.5 %), suggesting that strong regional demand (which drives region-wide surges) can spill over benefits even to areas that might not otherwise have as high demand. |

The seller-side estimates in Tables 8–12 offer a complementary perspective to the buyer-side results, and they illustrate the two-sided nature of the platform adjustments. First and foremost, higher pay clearly improves one outcome of great importance to sellers: their earnings per hour (Model S1). We see a positive coefficient of about 4 %, which means that on average a 10 % bump in the pay rate yields a 4 % increase in actual hourly earnings for drivers or couriers. This indicates that sellers do benefit from the surge, but not one-for-one

**Table 9: Estimated effects of pay increases on seller-side outcomes (Part 2)**

| Model | Outcome (unit) | $\beta$ : PriceChange $\gamma$ : | NeighborChange | Interpretation |
|---|---|---|---|---|
| S2 | Utilization (percentage points) | −1.0 (0.4) | −0.5 (0.2) | Sellers' utilization rates fall slightly when pay increases. A 10 % surge in pay corresponds to about a 1 percentage point decrease in the fraction of time that sellers are busy. This counter-intuitive result likely occurs because an attractive pay increase lures additional sellers into the market or into the area, increasing |

| | | | | |
|---|---|---|---|---|
| | | | | competition for each ride or task and leading to some idle time. In other words, supply overshoots demand in the short run. Additionally, if neighboring areas also have pay surges (indicating a broad region of high supply influx), local utilization drops a bit further (another half-point or so for full neighbor surge), consistent with a regional oversupply effect. |

**Table 10: Estimated effects of pay increases on seller-side outcomes (Part 3)**

| Model | Outcome (unit) | $\beta$ : PriceChange $\gamma$ : | NeighborChange | Interpretation |
|---|---|---|---|---|
| S3 | Completed Jobs (% change) | −2.0 (1.2) | −1.0 (0.6) | The total number of jobs completed by sellers in the focal area shows a slight decline when buyer prices increase (and thus demand falls). Even though more sellers might be available due to the pay incentive, the lack of commensurate demand growth means the total jobs per area or per seller may not increase; in fact, here we see a small negative change ( 2 % decrease). Neighboring surges also contribute to a reduction in completed jobs locally (about a 1 % drop if all neighbors have surges), per- |

| | | | | |
|---|---|---|---|---|
| | | | | haps because some demand is siphoned away or because local sellers shift to other areas. |

**Table 11: Estimated effects of pay increases on seller-side outcomes (Part 4)**

| Model | Outcome (unit) | $\beta$ : PriceChange $\gamma$ : | NeighborChange | Interpretation |
|---|---|---|---|---|
| S4 | Acceptance Rate (pp) | +3.0 (1.0) | +1.0 (0.5) | Sellers become more willing to accept incoming requests when pay is higher.<br><br>The acceptance rate rises by roughly 3 percentage points with a 10 % pay increase, indicating that drivers or couriers turn down fewer opportunities. This makes sense, as each job is now more profitable. There is also a mild positive spillover effect: if nearby areas also have high pay, acceptance rates locally tick up by about 1 point, possibly reflecting a generally high-demand environment or a contagion of "eagerness" to capitalize on peak times. |

— the less-than-proportional increase suggests that simply being paid more per job is partially offset by either doing slightly fewer jobs or incurring some additional non-earning time [22]. Indeed, looking at Model S2, we find that utilization drops by about 1 percentage point with a pay increase. This finding is somewhat counter-intuitive at first glance: one might expect that if drivers are earning more and presumably more motivated, they would spend more time actively engaged. However, the likely explanation is on the supply side: the surge in pay attracts additional sellers into the market or into the specific area, leading to more idle time for everyone until the market re-equilibrates. In essence, there are suddenly more drivers chasing each ride, so a given driver might end up waiting around a bit longer between rides. The negative γ on utilization (about -0.5) reinforces this interpretation — when the whole region is experiencing surges (meaning a widespread influx of drivers responding to high pay), the utilization dips a little further. This dynamic illustrates a classic

**Table 12: Estimated effects of pay increases on seller-side outcomes (Part 5)**

| Model | Outcome (unit) | $\beta$ : PriceChange $\gamma$ : | NeighborChange | Interpretation |
|---|---|---|---|---|
| S5 | Retention next week (pp) | +1.5 (0.7) | +0.5 (0.3) | Experiencing higher pay appears to slightly improve short-run retention of sellers. We find that the probability of a seller being active in the following week is about 1.5 percentage points higher if they experienced a surge pay this period. This suggests that some sellers are encouraged by the good earnings to continue working on the platform, at least in the near term. A smaller positive neighbor effect (around +0.5 pp if all neighbors surged) hints that a generally strong market in the region has a positive influence on participation. |

outcome in two-sided markets: incentives bring in supply quickly, which can overshoot momentary demand.

We also notice that the number of completed jobs (Model S3) does not increase and, in fact shows a slight decrease on average in our estimates, though this effect is modest. A coefficient of roughly -2% suggests that in the treated periods, perhaps due to the drop in demand from buyers facing higher prices, the total amount of work to be done actually shrinks a bit. Even if there are more drivers available, they cannot conjure up demand; and thus, if riders are taking fewer trips, the pie of available jobs may actually contract [23]. The spillover effect on jobs is similarly negative, at -1%, consistent with the idea that if nearby areas are also surging (and thus might be similarly seeing demand tempered by high prices), our focal area might not benefit from overflow jobs — it might even lose some as the whole region sees a synchronized ebb in requests.

On the more positive side for sellers, the acceptance rate of offers goes up notably under higher pay, or Model S4. We estimate about a 3 percentage point increase in acceptance when surge pay is in effect. What this suggests is that drivers/couriers become less picky and more willing to accept the rides or tasks that they are offered, presumably because each opportunity is now more lucrative. This can have some positive effects on marketplace efficiency; fewer requests go unfulfilled or bounce between drivers. Interestingly, there is a small positive neighbor effect as well (around 1 percentage point), which could mean that when it's busy and profitable everywhere (a city-wide surge scenario), drivers are in a mindset to work and may be generally more responsive to requests, or it could reflect that if one area is surging, a driver might decline a request to move to that area, but if all areas are surging, they simply take what they get.

Finally, we consider the longer-term (albeit short-horizon) outcome of seller retention in Model S5. The data suggest that sellers who experienced a high-paying period are slightly more likely to be active again in the following week, with a coefficient of about +1.5 percentage points [24]. This indicates a positive, though small, effect on retention: a good earning experience might encourage them to come back for more. It is important to note that this is a short-run effect; we did not analyze whether the effect persists beyond one week. The neighbor spillover for retention is also positive but even smaller (+0.5 pp), implying that a generally robust market environment contributes a bit to keeping sellers around, perhaps by fostering a sense that the platform is worthwhile.

As with the buyers, the spillover effects in the seller outcomes underscore that interventions do not happen in isolation. If surrounding areas are also offering high pay, our focal area's sellers feel some impact. In the case of earnings and acceptance, the neighbor effects are positive, which makes intuitive sense if we consider that a strong regional surge usually comes from strong regional demand — an environment where there are plenty of jobs to go around, even if a particular area's own demand is not sky-high. On the other hand, for utilization and jobs completed, the neighbor effects are negative, reinforcing the notion that an abundance of sellers region-wide can depress how much work each individual gets. These cross-effects, while smaller in magnitude than the direct impact of one's own treatment, are nonetheless important for understanding the full picture.

### 4.3 Interpreting Effect Sizes and Practical Significance

It is helpful to put the above estimates into perspective regarding their magnitude and practical significance [25]. The demand elasticity implied by the order volume response, around -0.6 for a 10 % price change, is moderate. This implies that the platform can modestly raise prices without losing all its customers, but a substantial drop in usage would likely outweigh the gain in per-order revenue if taken too far. In operational terms, seeing a 6 % decline in orders when prices go up by 10 % suggests a fairly responsive customer base. That kind of information is very valuable for pricing strategy: it quantifies the trade-off between higher revenue per order and fewer orders.

The wait time reduction, on the order of half a minute faster service for an isolated price increase, may be small in absolute terms. However, taking into account that typical wait times for these services could be just a few minutes, going down 0.5 minutes (30 seconds) could mean a quite noticeable improvement; think, for example, going from 5 minutes to 4.5 minutes. This could modestly improve customer satisfaction for those still ordering, partially compensating for the higher price in experience. On the other hand, if many areas are surging, loss of that benefit-as the neighbor effect adds back 0.2 minutes-is something the platform might consider in communication or mitigation strategies. [26]

The competitive effect of the +5 percentage points is not trivial: In a city with thousands of daily users, a one-digit percentage point increase in defections can amount to substantial numbers of users absolutely. This underlines how fast customers can react to price differences by multi-homing. Even a few percentage points of users switching can put pressure on the platform's market share, especially if it means those users may favor a competitor permanently after a bad experience. While the platform may gain higher margin on each transaction from those who stay, it risks sending some fraction of users to a rival-which is a serious consideration for long-term strategy. On the supply side, the 1 percentage point drop in utilization may not seem worrisome at first glance-after all, it's a small move on a base that might be, say, 50-60 % on average. Yet for driver experience, going from, say, 60 % to 59 % active time suggests slightly more waiting around, which can translate to frustration or wasted fuel. If drivers notice that during a surge they are not getting as many rides as they might expect, they might question the benefit of the surge. That hourly earnings still rise by 4 % on average provides reason to believe that drivers overall do earn more, but perhaps not as much more than they hoped given the headline pay rate increase. This difference between expected and realized earnings is an important nuance; it could be communicated to drivers in order to set realistic expectations, or it could be dealt with by throttling the inflow of new drivers during surges. [27]

This 3-percentage-point increase in acceptance rate is a meaningful improvement in marketplace efficiency. Suppose usually 85 %of requests are accepted, and it increases to 88 % that can reduce the number of unfilled requests and raise reliability for customers. What this shows is that drivers are responsive to incentives not just in showing up, but also in actively engaging much more with the presented opportunities. This responsiveness is just what surge pay is supposed to accomplish: make sure that when demand spikes, supply rises not only in quantity but also in readiness to serve. Lastly, the retention effect, modest though it is at 1.5 points, is significant for the platform's growth and stability of supply. In a huge workforce of gig-economy participants, a difference of one or two percentage points in retention could mean hundreds of extra active drivers in the following week for a big city. That does indicate that good short-term earnings might convince some fence-sitters to come back for at least another week of work. While we cannot say whether this effect persists or leads to long-term retention, it does imply a short-run boost in participation which could be crucial during times of peak demand [28]. In all, the magnitudes of the effects we estimate are mostly moderate but not trivial. They accord with practical expectations: users and providers both respond to monetary incentives in measurable ways. Importantly, none of the estimated effects are so large as to be implausible (for example, we do not see anything like a 50 % drop in demand or a 0.0 minute wait time), which lends credibility to our analysis. The effects also reveal the contours of the market's elasticity: demand is somewhat elastic but not hyper-sensitive; supply is highly responsive in quantity but that response dilutes individual gains; and competitive pressures are present but the majority of users still do not switch unless differences become very pronounced. These insights could help platform managers design pricing interventions that achieve desired outcomes without unintended side effects, and they highlight the value of considering not just direct impacts but spillovers and equilibrium adjustments as well.

# 5 Discussion

## 5.1 Direct Effects and Spillovers

**Direct effects**

Buyers: higher prices ⇒ orders ↓, cancels ↑, switching ↑.
Sellers: higher pay ⇒ online supply ↑, acceptance ↑, engagement ↑.

**Spillovers**

- Demand: price hikes next door dampen local demand (refrain/shift).
- Supply: pay boosts draw sellers from neighbors ⇒ local utilization can dip.
- Exposure: magnitudes scale with neighbor-share treated; effects are statistically and economically meaningful.

**Implications**

- Ignoring interference misstates impacts (treated areas leak effects; neighbors inherit them).
- Account for both direct and indirect effects in inference and forecasting.
- Patterns align with intuition: stress nearby ⇒ local demand ↓; attractive pay nearby ⇒ resource reallocation.

The estimated direct effects of pricing interventions presented for two-sided markets are intuitive from an economic perspective. Buyers facing higher prices do buy less, adhering to the law of demand [29]. This we saw from the reduction in order volumes, increased cancellations, and migration of customers to other platforms in the results. On the other hand, sellers respond to higher pay by coming online more and accepting more work, consistent with economic incentives increasing supply. These direct effects confirm that the marketplace participants are sensitive to price signals in the expected directions: buyers treat a higher price as a deterrent (leading them to cut back or seek alternatives), while sellers treat a higher reward as an encouragement (leading them to engage more).

Beyond these direct effects, our analysis shows clear evidence of spillover effects. Interference in this context is a situation where a price change in one locality or group affects the outcomes in other localities or groups, and we find clear evidence of such patterns. We found, for instance, that even in the absence of a price increase in their home area, an area-wide price shock can somewhat depress demand across neighboring areas-particularly because some consumers in the high-price area either forgo consumption completely or attempt to shift their consumption into the neighboring zones. In the same way, a high pay incentive in one area can siphon drivers from neighboring areas, affecting supply balance in those areas. This includes the slight dip in utilization observed from our results [30]. The spillovers we measured, such as the extra reduction in orders or additional increase in switching when many neighbors are also on price hikes, are statistically significant and of economic importance. They underscore the fact that two-sided marketplaces are connected structures: what happens in one node of the network-whether geographical or platform network-does not stay contained but ripples through user and provider behaviors. These results, from a platform perspective, reinforce the need to consider interference when analyzing any intervention. One can miss part of the overall impact of an intervention by looking at only the area or group directly treated because part of it leaks out to other areas; similarly, the other areas may exhibit changes that could be mistakenly attributed to their own conditions rather than spillovers from the treated group. Our study brings into sharp relief that both direct and indirect effects must be considered for a holistic understanding. Fortunately, in our setting, the direction of spillovers was consistent with what one would expect: generally, stress on one side in neighboring areas tended to cause parallel stress or shifts - high prices next door can dampen the local demand a bit - and attractive conditions next door can slightly improve local outcomes or siphon resources away. Understanding these dynamics can be useful in constructing better predictive models of market dynamics.

## 5.2 Two Unexpected Results

**Wait times fell under higher prices**

- Demand ↓ ⇒ shorter rider queue.
- If local supply stays roughly stable, matching is faster.
- System-wide surges/spillovers can offset this.

**Utilization dipped under higher pay**

- Pay ↑ ⇒ supply inflow > demand ⇒ idle time ↑.
- Rides/hour ↓ but pay/ride ↑ ⇒ earnings/hour ↑ (e.g., ≈+4%).

While most of the results followed predictable patterns, there are two findings that were somewhat unexpected at first glance and deserve further discussion. The first is the impact of buyer price increases on wait times. Intuitively, one might expect either nothing to happen to wait times or for them to get worse if the prices go up-for example, if drivers make less per ride due to fewer tips or something, they might be less motivated to accept quickly, etc. However, our analysis uncovered a slight decrease in wait times associated with higher prices. This makes intuitive sense in light of the basic supply-and-demand interaction: higher prices suppress demand substantially, so there are fewer riders competing for the attention of the drivers. Provided that the number of drivers does not fall off a cliff at the same time, the ones who are on the road can find riders more easily because the queue of waiting riders is shorter. In a word, when demand cools off, matching becomes more efficient and wait times decline [31]. This is a subtle equilibrium effect that might be missed if one only considered, say, driver incentives in isolation. It shows how a policy aimed at controlling demand-by raising prices-can have side benefits in terms of service quality (faster pickup) for the remaining users. We did see, however, that this improvement in wait times can be counteracted by broader conditions in the market: if all areas are surging, then drivers may reposition or riders may redistribute in ways that eliminate the local wait time advantage. So, the reduction in wait times is a nuanced outcome that appears under localized price increases but not necessarily under system-wide price pressure.

The second notable finding is that increasing pay for sellers can lead to a slight decrease in their utilization. Intuition might suggest that if drivers were being paid more, they would end up working more - either more drivers join and everyone stays busy, or the same drivers stay out longer and take more rides. What we found instead is a small dip in the fraction of time drivers are actively engaged. The likely reason is an overshoot in supply: a surge or bonus acts as a strong magnet for drivers - people who might not have driven at that time decide to log on, and drivers from other areas might drive into the surge zone. In the short run this influx of additional supply overshadows the increase in demand - since the demand side is simultaneously experiencing higher prices, which dampen requests [32]. Consequently, drivers spend a bit more time waiting for rides, because there are more of them relative to the ride requests available. In other words, while each individual ride pays more, those rides are a bit less frequent per driver. This resolves what might seem like a contradiction: how can earnings per hour go up while utilization goes down? The answer is that earnings per hour is a product of pay per ride and rides per hour; the pay per ride goes up significantly under surge, but the rides per hour - utilization - goes down slightly because of the supply glut, so the net is still an increase in earnings - as we saw, +4 %, but not as high as it could be if utilization had remained constant. This dynamic is an important insight for platform managers: simply throwing incentives to drivers will get them to come, but if too many come, the efficiency per driver drops. Ideally, a platform would want just enough extra drivers to meet the elevated demand, and no more, but in practice overshooting is common because drivers make decisions independently.

These two findings-that wait times are shorter under higher prices and utilization is lower under higher pay-highlight the complex and sometimes counterintuitive equilibration occurring in two-sided marketplaces. They constitute a reminder that such interventions have indirect effects that partially offset or alter the direct intent. Understanding these outcomes should help refine theories of how such markets operate and allow strategies to be adjusted to mitigate unwanted consequences-for example, perhaps tempering surge multipliers to avoid an oversupply of drivers that is too large. [33]

### 5.3 Practical Implications

**Interference-aware experiments**

- Evaluate price interventions with spillovers in mind.
- Use buffers or cluster randomization to contain leakage.
- Measure neighbor exposure (*NeighborChange*) and adjust estimates.

**Balance both sides**

- Price ↑: earnings ↑ but demand/satisfaction ↓.
- Price ↓: volume ↑ but supply participation may weaken.
- Monitor a suite of metrics; curb surge overshoot (caps/smoothing).

**Competition and geography**

- Users switch when prices bite; loyalty is fluid.
- Competitor surges create risks/opportunities; time promos accordingly.
- Spillovers across areas imply intra-platform area competition (streak bonuses, walk-to-save nudges).

These findings have a number of practical implications for platform management and the design of pricing policies or experiments [34]. First, our results strongly indicate that any price intervention should be evaluated with its spillovers in mind. In practical terms, this means it would be wise when running an experiment increasing prices in a set of cities or zones to include some buffer regions or control for neighbor effects in the analysis. Traditional A/B test designs, which assume no interference, may misestimate effects if users or drivers can move or switch between platforms. The approach is to randomize at a cluster level-for example, treat an entire city and compare with another city rather than patchwork areas within one city-to naturally contain spillovers within treated clusters. Another complementary approach is to measure the exposure of control units to treatment units, as we did with NeighborChange, and incorporate that into the analysis to adjust for interference. In short, accounting for interference, and designing with it in mind, results in more sound interpretations of experiments conducted on two-sided platforms.

Second, pricing decisions should balance the metrics of both sides of the market. A change put in place to improve one side's outcomes may be offset by effects on the other side. For example, increasing prices improves driver earnings-a supply-side positive-but, as we've seen, it will also reduce demand and may impair customer satisfaction or retention, a demand-side negative [35]. Similarly, lowering prices may please customers and increase order volume, but if it significantly cuts into driver earnings, the number of drivers participating or their effort may be reduced. Our study echoes that the health of the marketplace involves maintaining a balance between supply- and demand-side incentives. Accordingly, decisions such as introducing surge pricing or discounts need to track not just the target metric of revenue, demand, or supply engagement but a set of secondary metrics describing how the other side responds. In our example, the surge improved earnings but reduced utilization; a careful design might be to accompany the surge with measures to prevent too many drivers from crowding in-perhaps capping the number of drivers who can enter a zone, or smoothing out the surge multiplier-to preserve efficiency. The bottom line is that interventions must be designed and communicated with both sides of the market in mind, lest a problem simply be shifted from one side to the other.

Third, the competitive context always needs to be considered. We observed that a large fraction of users will switch to alternative services if they feel the price on one platform is too high. This suggests that platforms do not operate in a vacuum, and user loyalty can be relatively fluid in price-sensitive situations [36]. A practical implication is that platforms may coordinate promotions or at least pay attention to competitors' prices. If one platform increases prices (by choice or because the automated surge pricing kicks in), it could drive customers toward its competitors, which may have longer-term implications if the customers don't come back. On the other hand, when a platform recognizes a competitor is surging-say, due to a major event-it may exploit the situation by advertising its relatively lower price or keep its prices stable to capture the defectors. Another ramification of cross-geographic spillovers is that competition is not just platform-to-platform but also area-to-area within a platform. Users can and do move to where the deals are best. For example, one might walk a few blocks to avoid a surge zone for ride-hailing. In response, platforms have in some cases tried strategies such as offering "surge streak bonuses" to ensure drivers don't rush to chase surges elsewhere and that coverage is maintained, or inform customers that a short walk could yield a lower fare. All of these tactics revolve around the insights that users and providers are responsive and that they re-optimize their choices when prices change.

### 5.4 Identification, Assumptions, and Robustness

One fundamental assumption is that, conditional on controls and fixed effects we introduced, the timing and location of price changes are quasi-random from the perspective of the outcomes [37]. That is, we assume that there were no unobserved factors driving the occurrence of a surge - or promotion - and the jump in, say, order cancellations, over and above what we controlled for. This is a strong assumption, and in reality there's always the possibility that some confounder remains. For example, if the platform algorithmically triggers surge pricing in response to a sudden event - such as a big sports game ending - which independently also affects demand or supply in ways

we did not model, then our β captures not only the effect of the price change but also the effect of that event. To try to rule out such a scenario we included time fixed effects - which absorb any shock common to all areas in a given time - and introduced specific controls like holidays or weather that could precipitate both pricing changes and outcome shifts. We also excluded data from obviously special events which were unique to one area.
Another assumption is that interference is limited to a local neighborhood-so-called ”partial interference” in some causal inference literature. By including only the neighbors' treatment status, we implicitly assumed that units further away from the neighborhood have negligible spillovers. In reality, longer-range spillovers may exist-particularly in an area well-connected by transportation, or through social and information networks. If these longer-range spillovers do exist, our model risks underestimating the overall size of the spillover, or even misattributing some effects to incorrect sources [38]. However, we suspect that most adjustments by riders and drivers do occur quite locally-for instance, it's unlikely that a driver would drive 50 miles away just to chase a surge, and a customer is unlikely to factor in prices in some other city when considering ordering dinner. In that sense, the partial interference assumption seems reasonable within the scope of our study, but it is an assumption nonetheless that could be tested in future work-for instance, one could broaden the definition of neighbor and see whether additional terms contribute.
We performed several robustness checks to raise confidence in our results. First, we conducted placebo tests by assigning “fake” treatment periods before the actual price changes took place. Reassuringly, these placebo treatments did not yield significant effects on the outcomes, which supports the idea that we do not simply capture pre-existing trends. Second, we tried alternative definitions of the neighbor exposure. Instead of a simple share of adjacent areas, we weighted neighbors by their proximity or relative traffic flow to see whether that changed the inference; the results were qualitatively similar, indicating that our spillover detection is not overly sensitive to the exact way neighbors are counted. Third, we varied the threshold for what we considered a “price change” event-for example, using a 15 % increase criterion instead of 10 %, or looking at any non-zero surge multiplier rather than requiring a certain size [39]. The direction and significance of the main effects remained stable under these variations, though of course the magnitude of β would adjust somewhat as we made the treatment definition more or less strict. Fourth, we looked at subsets of the data, such as excluding major holidays or extreme weather days, and found our estimates persisted across these subsamples, which suggests the patterns we identified are general and not only tied to, say, holiday behaviors.
Though these checks temper the findings, we caution again that our study is observational and depends on the structure of the model for the interpretation of causal effects. There may be unobserved heterogeneity or dynamics we have not captured. For instance, drivers and riders learn and form expectations-if they believe prices are about to surge, they may adapt in advance, which would invalidate our exogeneity assumption. Also, our sources for data (surveys and public information) introduce measurement error: it is unlikely every price change event was correctly identified, and not every self-reported behavior change is accurate. These issues may attenuate some effects or add noise. We attempted to counteract errors by combining multiple signals and focusing on clear-cut cases of changes in prices, but some

uncertainty does remain. In all, we believe that our findings reflect real causal relationships in how pricing affects marketplace outcomes, bolstered by the consistency of results and robustness tests [40]. However, readers should interpret the exact magnitudes with appropriate caution. The main value of the analysis is to show the presence and direction of various effects-including some surprising or non-intuitive-rather than to pin down an exact universal constant of demand elasticity or supply response. Future work with more granular data, or experimental designs, might build on these insights and further validate the assumptions made here.

## 6 Conclusion

Pricing interventions in two-sided marketplaces create far-reaching chains of reactions on the part of both buyers and sellers, as well as along different parts of the market. In this paper, the authors combined survey data, public information, and time-based patterns to study these reactions in an observational setting. They have developed a simple OLS regression framework that estimates both the direct effect of a price (or pay) change on local participants and its indirect spillover effects on neighboring participants. This allows for a range of changes in outcomes to be quantified, from core metrics such as order volume and earnings to behavioral adjustments in application switching and service acceptance rates.

These findings paint a consistent picture of the supply and demand response to price changes in a two-sided market: as prices rise for buyers, demand moderates-fewer orders are issued, and more customers abandon requests or search for alternatives [41]. As demand falls, those customers who continue to use the service may see somewhat improved service (e.g., quicker wait times) as the system becomes less strained. As pay incentives rise, sellers work more hours and accept more work, but this increased supply can lead to diminishing returns in terms of how continuously busy each seller remains. This reflects the classic tango of supply and demand-price rises, demand falls; rewards rise, supply expands, sometimes overshooting equilibrium levels.

Importantly, the analysis shows that these effects spill over beyond the directly treated group. Spillovers presuppose that a change in price in one location or in one segment may affect the outcome in another location or segment. Thus, a surge pricing event in one location may have a spillover impact on demand and supply in other locations, and vice versa. The result points to the fact that platforms face network effects whereby policies in one segment may spill into another—sometimes with positive aggregate externalities, such as when higher demand trickles down to neighboring markets, but sometimes negative consequences emerge when user defection or oversupply decreases efficiency.

The methodological contribution lies in this use of available information in a particularly intuitive regression framework that serves as a valuable template for other applications in which randomized experimentation is either impossible or interference is significant. While this approach cannot replace controlled experiments, it offers a clear and believable method for estimating treatment and spillover effects in environments with observational data [42]. The estimates obtained are consistent in direction and plausible in magnitude, supporting the idea that meaningful causal insights about marketplace behavior can be extracted even from non-experimental data, provided confounding and interference are carefully addressed.

Upcoming research could extend this analysis in a variety of ways. First, it might explore heterogeneous effects: analyzing whether some neighborhoods or types of users are more responsive to changes in prices than others. Second, longer-term effects could be examined, since prolonged periods of high prices or higher pay may have implications for user retention and platform reputation in the longer run. A fuller analysis of dynamics across platforms would also be informative, integrating data from multiple competing services to capture substitution and competitive responses under surge or discount conditions. This study reinforces a central principle in the management of two-sided platforms: any intervention must be assessed in a broad systemwide framework. A price cannot be treated as an independent lever that acts only on a single metric or one group; it triggers an interdependent cascade of responses across a marketplace. Through measuring and acknowledging both first-order and spillover effects, decision-makers can better develop pricing algorithms, surge policies, and promotional strategies that effectively reach desired outcomes with minimal spillovers. The evidence and framework presented here contribute to a more nuanced understanding of marketplace dynamics and support the creation of analytical methods that explicitly consider-and incorporate-the interconnected nature of platform ecosystems. [43]

**References**

[1] X. Zhang, X. mei Yuan, and F. li Wei, "The dynamic broker dominant supply-demand matching decision using multiagent system,"
Mathematical Problems in Engineering, vol. 2020, pp. 1–11, 6 2020.

[2] A. Zrahia, N. Gandal, S. Markovich, and M. Riordan, "The simple economics of an external shock to a bug bounty platform," Journal of Cybersecurity, vol. 10, 1 2024.

[3] E. S. Oblander, S. Gupta, C. F. Mela, R. S. Winer, and D. R. Lehmann, "The past, present, and future of customer management,"
Marketing Letters, vol. 31, pp. 125–136, 6 2020.

[4] M. Haldar, M. Abdool, L. He, D. Davis, H. Gao, and S. Katariya, "Learning to rank diversely at airbnb," in Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, pp. 4609–4615, ACM, 10 2023.

[5] D. Pennino, M. Pizzonia, A. Vitaletti, and M. Zecchini, "Blockchain as iot economy enabler: A review of architectural aspects,"
Journal of Sensor and Actuator Networks, vol. 11, pp. 20–20, 3 2022.

[6] J. Konietzko, N. Bocken, and E. J. Hultink, Online Platforms and the Circular Economy, pp. 435–450. Springer International Publishing, 2 2019.

[7] S. Banerjee, S. Gollapudi, K. Kollias, and K. Munagala, "Www - segmenting two-sided markets," in Proceedings of the 26th Inter- national Conference on World Wide Web, pp. 63–72, International World Wide Web Conferences Steering Committee, 4 2017.

[8] A. N. Turi, J. Domingo-Ferrer, D. Sa´nchez, and D. Osmani, "A co-utility approach to the mesh economy: the crowd-based business model," Review of Managerial Science, vol. 11, pp. 411–442, 2 2016.

[9] F. Rowe and M. L. Markus, "Taking the measure of digital giants: Amazon and the social welfare computing research agenda,"
Electronic Markets, vol. 32, pp. 437–446, 4 2022.

[10] S. Gallo, "Fintech platforms: Lax or careful borrowers' screening?," Financial innovation, vol. 7, pp. 1–33, 7 2021.

[11] P. Mehendale and P. Buddiga, "Study of dynamic pricing model for two - sided

marketplace," International Journal of Science and Research (IJSR), vol. 10, pp. 1787–1793, 9 2021.

[12] K. Kang, B. Q. Tan, and R. Y. Zhong, "Cloud-based 3d printing service allocation models for mass customization," The International Journal of Advanced Manufacturing Technology, vol. 126, pp. 2129–2145, 3 2023.

[13] M. Haldar, H. Zhang, K. Bellare, S. Chen, S. Banerjee, X. Wang, M. Abdool, H. Gao, P. Tapadia, L. He, and S. Katariya, "Learning to rank for maps at airbnb," 6 2024.

[14] Q. Tang, Z. Shao, L. Huang, W. Yin, and Y. Dou, "Identifying influencing factors for data transactions: A case study from shanghai data exchange," Journal of Systems Science and Systems Engineering, vol. 29, pp. 697–708, 11 2020.

[15] F. Sarvi, M. Aliannejadi, S. Schelter, and M. de Rijke, "How to make an outlier? studying the effect of presentational features on the outlierness of items in product search results," in Proceedings of the 2023 Conference on Human Information Interaction and Retrieval, pp. 346–350, ACM, 3 2023.

[16] M. Peitz and P. Waelbroeck, "An economist's guide to digital music *," CESifo Economic Studies, vol. 51, pp. 359–428, 1 2005.

[17] P. Shi, "Optimal match recommendations in two-sided marketplaces with endogenous prices," SSRN Electronic Journal, 1 2022.

[18] V. Kumar, N. Nim, and A. Agarwal, "Platform-based mobile payments adoption in emerging and developed countries: Role of country-level heterogeneity and network effects," Journal of International Business Studies, vol. 52, pp. 1529–1558, 7 2020.

[19] P. Shi, "Optimal match recommendations in two-sided marketplaces with endogenous prices," in Proceedings of the 23rd ACM Conference on Economics and Computation, pp. 794–794, ACM, 7 2022.

[20] S. Bourke, "Recsys - the application of recommender systems in a multi site, multi domain environment," in Proceedings of the 9th ACM Conference on Recommender Systems, pp. 229–229, ACM, 9 2015.

[21] T. Widmer, P. Karaenke, and V. Sugumaran, "Two-sided service markets: Effects of quality differentiation on market efficiency,"
Managerial and Decision Economics, vol. 42, pp. 588–604, 11 2020.

[22] A. Aldweesh, "Blockticket: A framework for electronic tickets based on smart contract.," PloS one, vol. 18, pp. e0284166–e0284166, 4 2023.

[23] D. Wigan, Combating Fiscal Fraud and Empowering Regulators - Uber Global Wealth Chains, pp. 194–214. Oxford University PressOxford, 2 2021.

[24] M. Grbovic and T. Noulas, "Wsdm - workshop on two-sided marketplace optimization: Search,pricing, matching & growth," in
Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, pp. 801–802, ACM, 2 2018.

[25] B. Shi, E. H. Gerding, P. Vytelingum, and N. R. Jennings, "An equilibrium analysis of market selection strategies and fee strategies in competing double auction marketplaces," Autonomous Agents and Multi-Agent Systems, vol. 26, pp. 245–287, 1 2012.

[26] M. Haldar, H. Zhang, K. Bellare, S. Chen, S. Banerjee, X. Wang, M. Abdool, H. Gao, P. Tapadia, L. He, and S. Katariya, "Learning to rank for maps at airbnb," in Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining,

pp. 5061–5069, ACM, 8 2024.

[27] T. M. Key and G. Challagalla, "Information transparency as competitive advantage: research approaches to platform ecosystem design and competition," AMS Review, vol. 10, pp. 145–147, 4 2020.

[28] X. Tang, H. Lu, W. Huang, and S. Liu, "Investment decisions and pricing strategies of crowdfunding players: In a two-sided crowdfunding market," Electronic Commerce Research, vol. 23, pp. 1209–1240, 9 2021.

[29] null AlijaniReza, null BanerjeeSiddhartha, null GollapudiSreenivas, null KolliasKostas, and null MunagalaKamesh, "The segmentation-thickness tradeoff in online marketplaces," Proceedings of the ACM on Measurement and Analysis of Computing Systems, vol. 3, pp. 1–26, 3 2019.

[30] Y. Feng, R. Niazadeh, and A. Saberi, "Two-stage stochastic matching and pricing with applications to ride hailing," 1 2022.

[31] G. Kamulegeya, R. Mugwanya, and R. Hebig, "Measurements in the early stage software start-ups: A multiple case study in a nascent ecosystem," Foundations of Computing and Decision Sciences, vol. 43, pp. 251–280, 12 2018.

[32] E. Fang, X. Li, M. Huang, and R. W. Palmatier, "Direct and indirect effects of buyers and sellers on search advertising revenues in business-to-business electronic platforms," Journal of Marketing Research, vol. 52, pp. 407–422, 6 2015.

[33] M. Bhattacharya and A. Barapatre, "Recsys - query as context for item-to-item recommendation," in Fourteenth ACM Conference on Recommender Systems, pp. 575–576, ACM, 9 2020.

[34] G. Stein, "Social cost," Gabe Stein, 7 2022.

[35] R. Bala, K. R. Sarangee, S. He, and G. Jin, "Get us ppe: A self-organizing platform ecosystem for supply chain optimization during covid-19," Sustainability, vol. 14, pp. 3175–3175, 3 2022.

[36] L. Yuan, C. Gao, A. Leung, and Q. Ye, "Unraveling information asymmetry in blockchain-enabled nft marketplaces: The impact of rarity rank on consumer behavior," 1 2024.

[37] H. Zhang, H. Wang, A. Mehra, and E. Zheng, "Chasing market growth and matching efficiency in two-sided platforms: Evidence from the lazy-minting policy in an nft marketplace," SSRN Electronic Journal, 1 2022.

[38] M. Drouillard, Addressing Voids: How Digital Start-ups in Kenya Create Market Infrastructure, pp. 97–131. Palgrave Macmillan UK, 11 2016.

[39] A. de Ruijter, O. Cats, R. Kucharski, and H. van Lint, "Evolution of labour supply in ridesourcing," Transportmetrica B: Transport Dynamics, vol. 10, pp. 599–626, 1 2022.

[40] A. Chin and Z. T. Qin, "A unified representation framework for rideshare marketplace equilibrium and efficiency," in Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems, pp. 1–11, ACM, 11 2023.

[41] G. Dushnitsky and T. Klueter, "Is there an ebay for ideas? insights from online knowledge marketplaces," European Management Review, vol. 8, pp. 17–32, 2 2011.

[42] S. S. S. Huda, A. Akhtar, I. Uddin, S. R. Maliha, and S. M. Mowla, ""truck lagbe" app powered goods transportation," Journal of Information Technology Teaching Cases, vol. 14, pp. 18–23, 12 2022.

[43] K. S. Micken, S. D. Roberts, and J. D. Oliver, "The digital continuum: the influence of

ownership, access, control, and cocreation on digital offerings," AMS Review, vol. 10, pp. 98–115, 7 2019.