Pre-Analysis Plan: Price Information and Competitive Spillovers in an Online Platform in Pakistan

Shotaro Nakamura, Syed Ali Hasanain, and Adeel Tariq

June 13, 2021

Abstract

Information and search frictions in developing markets may induce agents to make distorted pricing and advertising choices, which in turn may generate further negative externalities that reinforce such frictions. While direct effects of information interventions to alleviate some of these frictions are extensively studied (e.g. Aker 2010; Aker and Mbiti 2010; Jensen 2007), potential spillovers and their mechanisms remain relatively under-explored. We wish to provide empirical insights into the ways in which new, individual price information in a developing economy alters market conditions, or ways in which the effects of new information are altered by such market conditions. To this end, we will conduct a private price information intervention to sellers on a listing platform for used vehicles in Pakistan, PakWheels.com, from June to August 2021. We will generate variation in the timing of treatment by market sub-sections via a two-step cluster-randomization design, in order to separate direct treatment effects from spillovers. In our primary analysis, we will measure direct and spillover effects on a) changes to the listing price, b) occurrence of transaction, c) transaction price, d) an advertising index, and e) page views as a proxy for buyer interest. In our secondary analysis, we will identify ways in which the treatment effects interact with, or in turn affect, market efficiency and structure.

Contents

1 Introduction

Information friction plays a critical role in the efficiency of emerging markets, which often see high price levels and dispersion (e.g. Allen 2014; Andrabi et al. 2017). Interventions that leverage on information communications technology (ICT) are shown to reduce such information friction (e.g. Aker 2010; Aker and Mbiti 2010; Jensen 2007), yet the mechanisms through which such interventions trigger strategic responses, spillovers, and adjustments to market structure remain underexplored. On one hand, it may be possible that large enough exogenous shocks in information friction may induce general equilibrium shifts, for instance by improving efficiency of the supply channel (e.g. Jensen and Miller 2018). On the other hand, the market structure may determine how participants internalize information friction, and limit the extent to which interventions could shift market equilibria (Mitra et al. 2018).

Our aim is to improve our understanding on links between information friction, individual choices, and spillovers in a developing-market setting. We study the mechanisms by which information frictions may alter agents' strategic behaviors in the market, generate spillover effects onto their competitors, and affect the competitive landscape. We are running a randomized controlled trials (RCTs) in which we provide price information—called the Price Calculator—privately to sellers on a listing platform for used vehicles in Pakistan, PakWheels.com. We generate variation in the timing of treatment by market sub-sections via a two-step cluster-randomization design, in order to separate direct treatment effects from spillovers.

It has traditionally been difficult to measure spillovers and strategic responses to interventions in developing-market settings, due to logistical constraints and the lack of data. The increasing popularity of online platforms and marketplaces in developing countries, however, presents us with an opportunity to study market behaviors more extensively and comprehensively than before. Unique data captured by our implementation partner PakWheels allow us to measure changes in sellers' strategic choices (pricing and advertising) and demand-side responses (page views and inquiries). We are also able to trace spillover effects and other market-level consequences of the intervention, because we observe online activities for the full universe of listings on this platform. We are also able to infer potential channels of such spillovers, as we observe not only sellers' choices on multiple strategic dimensions, but also of shifts in buyer attention and on the quantity of vehicles on the market. Furthermore, by providing the Price Calculator estimate privately we can rule out direct information spillovers and instead focusing on knock-on effects of *choices* that treated sellers make.

This project highlights the potential for online platforms in emerging markets to address economically important questions and help reduce frictions and constraints. Our partner, PakWheels.com, is the leading and nationally recognized online classified advertisement platform for used vehicles Pakistan, allowing us to make a stronger claim to external validity than other information interventions with smaller and selected subsamples. Our intervention are also being conducted at low cost and within the framework of existing online platform, allowing us to assess, for instance, the potential efficacy of online-platform based interventions to improve small to medium enterprises' (SME) performance. E-commerce in emerging markets is growing in importance, and there is burgeoning evidence on how online markets help improve access to markets for those who were traditionally excluded (Couture et al. 2018). Findings from our research may have implications on how online platforms could help level information asymmetry against individual sellers and small used-car dealers by providing them with market insights. If proven effective, this study may serve as an example of cost-effective, platform-based services to improve SME performance, which would be particularly interesting in light of low cost-effectiveness of business training interventions (McKenzie and Woodruff 2014; Blattman and Ralston 2015).

The remainder of this pre-analysis plan is organized as follows; Section 2 highlights the primary and secondary research questions we intend to answer with our intervention. Section 3 provides a summary of existing empirical evidence on search and information frictions in developing economies, insights from research on digital platforms, and the emergence of new literature on platform experiments in such contexts. Section 4 presents a summary of predictions from the theoretical framework, and highlights connections to the research questions in Section 2. Section 5 describes the research design and results of power analysis. Sections 6 and 7 discuss data sources and outcomes, respectively. Section 8 pre-specifies the analysis we intend to run, and draws connections with the theoretical framework and the research questions.

2 Research questions

Our empirical objective is to understand how an information intervention induces changes in pricing, affects sellers' market outcomes, and generates spillovers that may have implications on the market structure. We divide our research questions into the following two strands.

The first strand of questions focuses on direct treatment effects on listing prices, transaction outcomes, and mechanisms at the individual level. Our overall hypothesis is that the information intervention reduces noise in sellers' beliefs about the distribution of demand, and improves their market outcomes. We hypothesize *a priori* that there may be some countervailing mechanisms that interact with the noise in beliefs, and that such mechanisms may be triggered as a result of the intervention. Our primary empirical goals are to measure average direct treatment effects on pricing and transaction outcomes, and to identify the relevance of any countervailing mechanisms.

The second strand of questions concerns spillovers and other market-level impact of the information intervention. We focus on the following potential channels: a) diffusion of information itself, b) competitive responses to treated individuals' strategic choices, and c) reduction in search frictions and congestion in the market. Our primary empirical goal is to identify if our intervention induces spillovers and demand-side responses. Our secondary empirical goal is then to identify channels of the spillover, such as ones suggested above.

Following is the list of primary (in bold) and secondary questions, with links to the theoretical predictions in Section 4.

1. Does the price information intervention induce individual responses?
   - 1.1. **Does it induce sellers to adjust their listing prices toward recommended prices?** (Prediction 1.)
     - 1.1.1. Does it induce level shifts in prices?
     - 1.1.2. Does it induce sellers to update their listing prices after the initial listing?
     - 1.1.3. Across what characteristics do we observe heterogeneous treatment effects?
       - sellers' experience
       - product heterogeneity in market segments
       - asymmetrical response above/below suggested price
   - 1.2. **Does the price information intervention improve sellers' returns from listing on the platform?** (Prediction 2.)
     - 1.2.1. **Does it increase the probability that treated vehicles are sold?**
     - 1.2.2. **And/or does it affect the transaction price?**
   - 1.3. Does the intervention affect sellers' beliefs about "expected" transaction prices?
   - 1.4. **Do sellers respond to the intervention by making other strategic adjustments, such as advertising?** (Prediction 3.)
2. Does the price information intervention create spillovers and other knock-on effects?
   - 2.1. **Do untreated sellers adjust their listing prices when their peers in the same market cluster are treated?** (Prediction 4.)
   - 2.2. Are spillover effect larger for individuals and market segments with higher pre-intervention information friction?
     - 2.2.1. Is the effect larger for inexperienced sellers than for experienced ones?
     - 2.2.2. Is the effect larger in markets with more product heterogeneity?
     - 2.2.3. Is the effect larger in markets in which price information was limited prior to the intervention?
   - 2.3. **Does the intervention induce a shift in buyer attention toward treated sellers?**
   - 2.4. Does the intervention affect the duration of posts' life and the level of congestion in treated market segments? (Prediction 4.)

3 Literature Review

In this literature review, we discuss the empirical background to our research questions and provided a justification for our focus on online platforms in developing economies. We first briefly reiterate the motivation of our research around information and search frictions in developing markets, as we already discussed in Section 1. We then present empirical evidence from online platforms in developed economies, and discuss their role in reducing search and information frictions as well as their limitations. We conclude by highlighting an emerging body of evidence from online-platform interventions in developing countries and identifying our potential contributions.

3.1 Search and information frictions in developing markets

High price levels and dispersion have often been cited as a result of search frictions, trade costs, and market power in developing markets and supply chains (e.g. Allen 2014; Atkin and Donaldson 2015). The roles of innovations in information communications technology (ICT) in reducing such frictions and market failures have been extensively studied (e.g. Aker 2010; Aker and Mbiti 2010;Andrabi et al. 2017; Jensen 2007). Yet, due in part to the dearth of experiments with variations at the market-segment level, we know relatively little about market-level implications of such interventions.

It is possible that large enough exogenous shocks to information friction induce knock-on effects on efficiency, up through supply channels or in other general-equilibrium sense (e.g. Jensen and Miller 2018, Hasanain et al. 2019). Yet our insights are limited on externalities generated *within markets* by information interventions, particularly on how they affect strategic choices of individuals and their competitors alike, and implications on market efficiency. Jensen (2007) and Aker (2010), for instance, are primarily focused on showing convergence in commodity prices in market-segment-wide interventions that make it difficult to capture individual mechanisms or spillover effects. Other subsequent experimental work on price information, on the other hand, often focus on small selected subsample and lack the sample or the design to speak to spillovers and market-wide effects.

Of the papers that speaks to spillovers and market-wide effects, one finds that market structure and form of negotiation may determine how agents respond to information and search frictions. Mitra et al. (2018), for instance, shows that market provision of price information to potato farmers does not affect farm-gate prices and revenue on average, but increases pass-through from middlemen to farmers. This implies that information affects the bargaining power of farmers, but overall effects on price is contingent upon the market structure in which they operate. There is also an emerging body of work suggesting the importance of relational contracting as a mechanism of coping with search and contracting frictions (e.g. Startz 2018). The main tenant of our work is similar to that on relational contracting, in that we hope to understand the roles of mechanisms that interact with information friction, such as agents' beliefs about demand and their strategic choices. We then hope to identify the existence and channels of externalities stemming from those individual beliefs and strategies, and how information interventions would interact with—or in turn affect—the market structure.

3.2 Roles of online platforms

High levels of information and search frictions make developing-economy markets a particularly appropriate context to pursue interventions based on online platforms, particularly on the tail of previous experiences with other ICT interventions. The premise of many online platforms is that they make it much easier to acquire information about products on the market and provide other services that reduces search and time costs so as to make the experience as seemless as possible. Yet, evidence of persistent price dispersion in **developed** economies' online markets suggests that platforms will not completely eliminate information and search frictions in developing economies either (Einav et al. 2015; Horton 2019; Fradkin 2015). The question is why search and information frictions persist in a world with plausibly low search and information costs.

One view is that price dispersion and frictions on online platforms are in part endogenous choices that platform operators make relative to the competitive pressure they want to induce. The following tension is at play; on one hand, platforms may want to reduce search friction and guide buyers to a small set of products that match their preferences. On the other hand, platforms also hope to induce price competition between sellers, which becomes more challenging when markets become highly segmented by specific search criteria. This idea is explored by Dinerstein et al. (2018), who evaluate the impact of a redesign by eBay that directed consumers towards products they prefer while inducing stronger price competition among sellers. Estimated search frictions and online retail margins suggest that this particular redesign in the search process reduced price levels and variation, suggesting that balance between low information friction and competitive pressure is key to efficient online markets. This trade-off may be even more salient in developing economies with higher existing frictions and other market failures.

3.3 Platform interventions in developing markets

Given the potential trade-off between search frictions and competitive pressure, what happens when online platforms in developing economies are introduced or redesigned to reduce search and information frictions? An emerging body of evidence from the literature on online platforms in developing countries provide some insights into this question, albeit with limitations. We summarize the evidence and highlight our potential contributions.

First, it seems that online platforms can reduce information and search frictions and improve welfare, although their successes may be varied. Couture et al. (2018), for example, show that while the benefits of access to e-commerce for rural markets in China are sizable, most of the gains accrue to the consumption side and to a minority of younger and richer users. The findings suggest that simply increasing access does not induce investments required to drive adaptation to e-commerce. In online labor market platforms, studies like Fernando et al. (2020) and Jeong (2020) found positive employment gains for disadvantaged groups, and reductions in wage price dispersion by reallocation of labor. The existence of heterogeneous effects therefore suggests the importance of understanding the mechanisms at play.

Second, there is evidence that search and information frictions still play a major source of inefficiency on online platforms in developing countries. Bai et al. (2020) describe the existence of such frictions on a Chinese platform AliExpress, and shows that positive shocks to demand and information improve firms' performance in the long run, independent of productivity or quality. This suggests that market dynamics may generate inefficient firms and low-quality goods to persist in markets with information and search frictions, based on the "luck" of having received positive initial demand shocks.

These strands of work point to potential benefits and limitations of online platforms in reducing information frictions, with welfare implications. Gaps in the literature still remain on a) how market participants internalize, and compensate for information and other frictions, and b) spillovers and systematic implications of externally adjusted information environment. We hope to contribute on these points by i) evaluating sellers' responses on a wide range of behaviors on the platform, and ii) systematically measuring spillovers by conducting a platform intervention that induces a shift in the information set at market segment levels of a dominant online platform in a key sector in Pakistan.

4 Theoretical Framework

We present a simple conceptual framework that describes a search process in which sellers set listing prices and place advertisements under noisy beliefs about demand. We derive predictions on how a price-information intervention may reduces noise in sellers' beliefs and affect list pricing, advertising, and market outcomes. The model is a simple static search framework based on Stigler (1961) and Diamond (1982). Most canonical models that focus on information assume that agents have full knowledge of factors like market friction and demand distributions (as summarized in Baye et al. 2007, as well as in papers like Allen 2014). We depart from this standard setup as follows:

- Sellers have biased or noisy beliefs about the distribution of buyer willingness-to-pay (WTP).
- This, along with noisy beliefs about the match rate and efficacy of advertising, may lead to biased or noisy beliefs about the probability of sale and to suboptimal pricing.
- Sellers can influence the match rate with potential buyers by engaging in costly actions, i.e. advertising.

We set up a model in which a seller i is endowed with an asset—a used vehicle—with characteristics s_i and information set I_i. Search and transaction occur through the following steps:

- Seller i forms a belief about the distribution of buyers' WTP, based on information I_i.
- Seller i chooses a listing price p_i^l and amount of advertisements a to optimize expected returns from participating in the marketplace.
- Seller i matches with a potential buyer via a Poisson process.
- Once matched, seller i makes a take-it-or-leave-it (TIOLI) offer p_i^t below p_i^l to the potential buyer.
- Transaction occurs if the matched buyer's WTP is higher than p_i^t.

We provide details on the set-up and derive the model in Appendix A.

4.1 Theoretical predictions

We derive the following predictions from the theoretical framework, and connect them to our main research questions in Section 2:

1. The price information intervention brings the listing price p_i^l closer to what it would be under no noise in beliefs about demand. (Research question 1.1.)
2. The information intervention increases expected returns from the search process. (Research question 1.2.)
3. The information intervention increases the consumption of advertising a if sellers' beliefs about expected returns from search are adjusted upward. (Research question 1.4.)
4. Spillover effects could occur through reduced noise in beliefs about demand, changes in the distribution of pageviews within or across market clusters, or changes in the Poisson match rate. (Research questions 2.1., 2.2., and 2.4.)

5 Research Design

We will conduct a field experiment in which we privately provide Price Calculator estimates to a randomly chosen subset of sellers. The Price Calculator estimates are based on a machine learning model using data on self-reported transaction prices from previous listings collected by PakWheels. The experiment is conducted within PakWheels' web and Android platforms, when sellers create new listings (i.e. the "Post an Ad" sequence). We assign treatment via a blocked, two-step randomization procedure, as described in Section 5.2. This intervention will be rolled out over the course of 8 weeks.

5.1 Sample selection

As of June 2021, the platform receives upward of 100,000 valid listings per month. Our sample is new posts on the platform during intervention period, except those for which PakWheels do not have sufficient data points to provide a Price Calculator estimate. The exact criteria for inclusion into the sample are masked for confidentiality reasons, but we expect to include approximately 88% of all new posts into the study sample, consisting of 73 distinct make-models.

5.2 Two-step treatment assignment procedure

Our two-step randomization process is as follows. In step 1, we block-randomize vehicle clusters, defined as the make-model (e.g. Toyota Corolla), into 9 treatment groups. These group numbers correspond to the week in which treatment starts being given to some in the cluster. In step 2, we randomize posts into treatment based on the last digit of the user ID on PakWheels. In order to ensure that treatment and control groups are comparable over the primary outcome variables, we test for balance using listings data from a pre-treatment period with the same sample inclusion criteria and randomization procedure as the experiment. We bootstrap-sample and iterate the randomization procedure 500 times, and identify seeds for which we fail to reject differences in all primary outcome variables (described in Section 7), adjusted for the false discover rate, at 5%. We then randomly select one of the qualifying seeds.

In step 1, we block make-model clusters over standardized cluster-level means of:

- log(absolute difference between listing price and Price Calculator estimate)
- 1 if reported as sold
- log(self-reported transaction price)
- advertising index[1]
- buyer-attention index
- cluster size.

The variables above, other than the cluster size, are primary outcome measures. Detail on how they are constructed is given in Section 7. Blocking is done with R's *blockTools* package (Moore 2012), which uses the optimal-greedy algorithm over the Mahalambois distance. We weight the five main outcome variables twice as heavily as the cluster size variable. Our

[1]There is a minor difference in definitions of constituent variables, due to limitations in the pre-intervention data

choice of weights is admittedly arbitrary, but the rationale is that the primary objective is to balance over main outcome variables, and then with cluster size.

Using these groupings, we identify first-stage assignments, i.e. the week in which some listings in a given cluster start receiving treatment. This assignment process is illustrated in Table 1. In "Treat" clusters, we randomly select 50% of new listings to receive the Price Calculator estimates, while none of the new listings in the "Control" clusters receive them. On the 9th week, all new listings (i.e. "Everyone") receive the Price Calculator.

Table 1: Cluster assignment schedule over time periods

Assignment	**Week 1**	**Week 2**	**Week 3**	...	**Week 8**	**Week 9**
1	Treat	Treat	Treat	...	Treat	Everyone
2	Control	Treat	Treat	...	Treat	Everyone
3	Control	Control	Treat	...	Treat	Everyone
...	...	...	...	...	...	...
8	Control	Control	Control	...	Treat	Everyone
9	Control	Control	Control	...	Control	Everyone

In step 2, post-level randomization is done based on the last digit of sellers' user-ID on PakWheels.[2] The choice of digits for treatment is fixed across cluster and time, in order to limit the extent of potential interference and for logistical simplicity. In other words, if a seller with user-ID i is in a treatment group for model m in week w, then all other future posts by i in m will be treated, as well as any other treated model m' in weeks $\geq w$. These digits are chosen by a random number generator.

5.2.1 Spillovers between clusters

One potential empirical challenges is interference between assignment clusters at the first stage of the randomization procedure. The concern is that if we define clusters too narrowly and pricing or advertising choices in one cluster could affect those in another, we would violate the Stable Unit Treatment Value Assumption (SUTVA). We intend to allay this concern by using a relatively broad definition of clusters as the make-model. We base our decision on observations from aggregated search logs data, details on which are provided in Section 6.5.1. We then address possible ways in which interference across clusters could still occur and their magnitudes.

First, we observe that a majority (58%) of specified searches included the make-model, and the majority of those 58% also included additional terms (e.g. model year, city, price ranges). On the other hand, 32 percent of specified searches did not include make-models,

[2]The reason for this randomization procedure, as opposed to some other procedures that does not rely on the user-ID, is because we are assigning treatment to a *flow* of new listings (and some new users), meaning that we can cannot necessarily make treatment assignments to a sample of known posts.

but instead included other types of qualifiers (such as city location, vehicle make only, or if the ad was featured or had pictures), or no qualifiers at all. These searches tend to be fairly broad, which led us to believe that they are primarily speculative searches that are unlikely to lead to meaningful actions by potential buyers. We do not have information (due to the capacity constraint at the firm for our data requests) on the remaining 10 percent of less frequent combinations of specified search terms. Overall, the breakdown of specified searches indicates that the make-model is likely a reasonable, and perhaps conservative, level of clustering, and any finer level of grouping could have meaningful interference between clusters.

Second, we believe that any interference across make-model clusters is likely minor given our experimental setup. This is because we provide private information that is specific to the posts' characteristic, which makes it unlikely that there would be large direct information spillover effect from a treated make-model to untreated ones. This still leaves room for interference between posts in the same cluster, but this is an effect we want to capture as spillovers.

Yet, the following are some of the ways in which interference *across* make-models, which would violate SUTVA in our context, could occur:

- Listing prices of one make-model is informative to sellers of another. If the intervention significantly altered the distribution of listing prices in treated make-model clusters, then spillovers of price information would occur to sellers in control clusters.
- Changes to the listing prices or advertising in treated clusters may shift buyers' attention to/from untreated make-models. Changes in buyer attention in untreated make-models may affect sellers' pricing and advertising choices.

We plan on running robustness checks to address these concerns. First we can include moments of listing prices of similar make-models as controls in the main estimating equation. Second, we can include the treatment status of similar make-models to capture spillovers beyond their own make-model clusters. We will define "similar models" as the randomization blocks identified in 5.2, as well as via an attributes-based approach (e.g. mid-range sedan, domestically manufactured *kei* cars). Third, we can empirically test if page views shift from untreated to treated make-models.

5.3 Treatment assignment and take-up

The intervention is designed so as to minimize non-compliance; those randomly assigned treatment are automatically shown the Price Calculator estimate on the interface while they create a post. One exception is the potential for selection into treatment based on users' app usage; users of PakWheels' mobile app would only be able to see the Price Calculator estimates (if assigned to treatment) once they update the app after the beginning of the intervention period. In contrast, users of PakWheels.com on the web (including internet browsers on mobile phones) would automatically be show the Price Calculator estimate, if they are assigned to treatment. This may generate selection into treatment based on a) users' preference for PakWheels' app and b) their propensity to update the app. We cannot derive conditional treatment effects, as we do not observe if an individual logged in using the app or web, or if their app is updated. To this end, we plan on identifying both intend-to-treat and treatment-on-treated effects, as highlighted in Sections 8.2.2 and 8.2.3.

5.4 Intervention instrument: The Price Calculator

We will provide estimates of the transaction price for used vehicles on PakWheels while sellers are creating their posts. The price information, which PakWheels calls "the Price Calculator", is based on a machine learning model trained to predict self-reported transaction prices using the firm's database on previous listings, conditional on a) the occurrence of transaction and b) observable attributes of the vehicle, but not of sellers' characteristics. Our hypothesis is that this information would help sellers identify realistic transaction prices, and set listing prices accordingly.

To identify an error-minimizing forecast model, we take a gradient boosting approach primarily for two reasons. First, gradient boosting—a method of ensemble predictions of many tree-based models—would allow us to identify a predictive model that is not confined to the make-model-modelyear categorizations. This was beneficial, as it allowed us to predict transaction prices for vehicles that had relatively small number of observations within their own make-model-modelyear, but for which we had sufficient information to provide predictions. Second is that the gradient boosting approach performed best in most measures of error against other approaches in our initial design process. This is consistent with the success of gradient boosting models in recent prediction competitions.

5.4.1 Display of the Price Calculator estimate

On PakWheels' web platform and mobile apps, sellers can create a new post by clicking on "Post an Ad." They are first asked to log in, so that we can identify the user-ID for each post. Users would not know their own user-ID (it is internal to PakWheels) or for which last digits we are providing Price Calculator estimates. Once logged in, users are asked to provide information about the vehicle they intend to sell, as shown in Figure 1. They then set the listing price in a box shown in Figure 2. If the seller is assigned to treatment, they are then shown a Price Calculator estimate, i.e. the machine-learning based transaction price forecast, as well as the 10th and 90th percentiles of reported transaction prices for the make-model-model year (or MMMY-version for frequently traded models). Figure 3 shows how the Price Calculator estimate is displayed along with a brief description. Treated sellers are then given a chance to update their listing price.

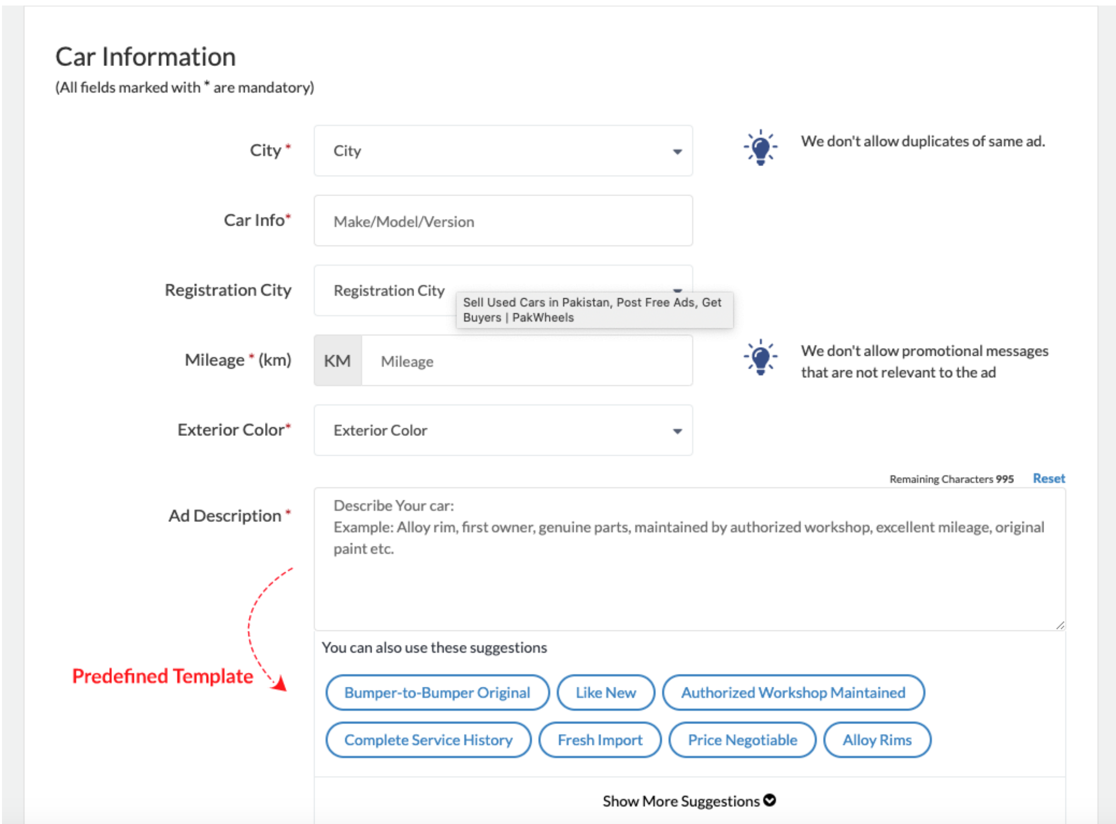

Figure 1: Making of a listing: Vehicle information

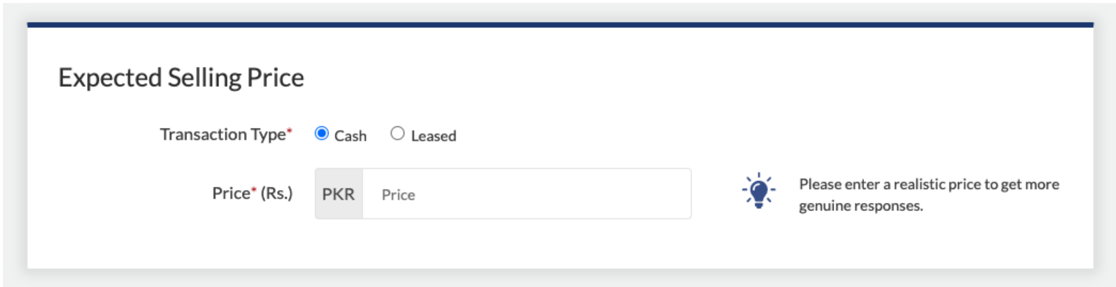

Figure 2: Making of a listing: Vehicle price

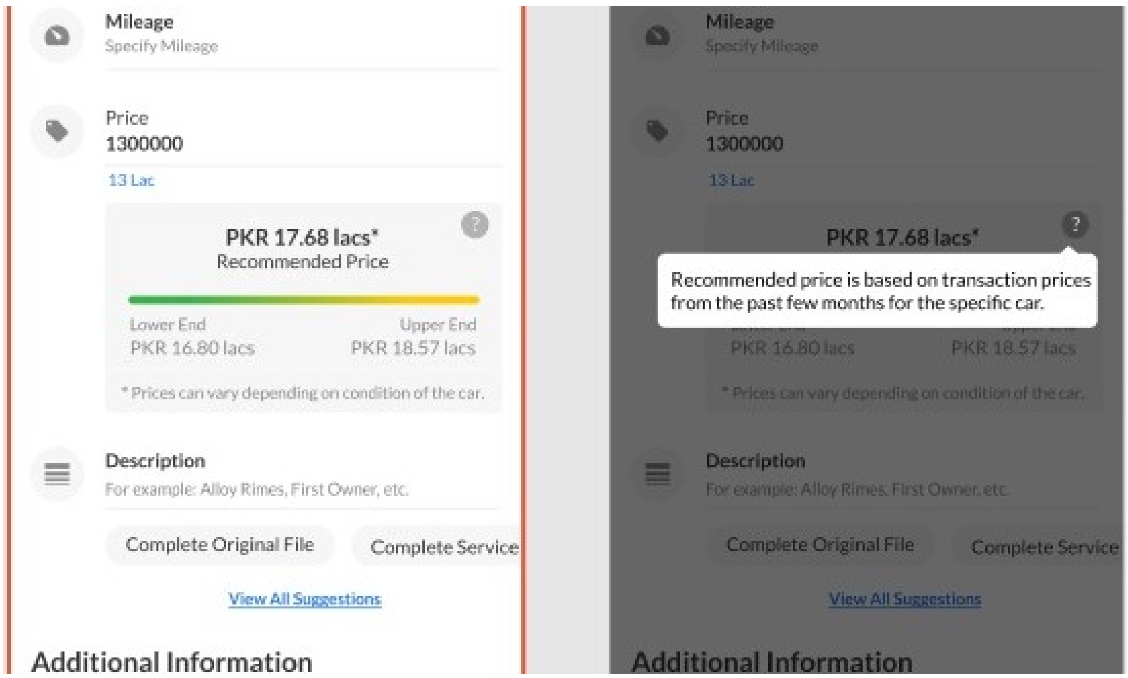

Figure 3: Display of the Price Calculator estimate

5.5 Unit of analysis

We will use the individual post as the unit of analysis for all pre-specified primary outcomes. This is because the treatment is provided at the post level, and we measure the spillover effect across posts. We also plan on conducting analysis at different units for some secondary outcomes and robustness checks. For instance, we will run analysis at the seller level as robustness checks to allay the concern that within-seller spillover would violate SUTVA. We will also run analysis at the model-week level on a range of secondary outcomes, such as the average duration of posts on the platform, as well as higher-order moments of the primary outcome variables (prices, buyer-attention index). Further details on primary and secondary analyses are provided in Section 8.

5.6 Statistical Power

Through simulations using actual data from PakWheels, we estimate the statistical power of detecting a range of effect sizes for both direct impact and spillovers. The data are from posts created between December 2020 and March, 2021 and closed by the time of power analysis (April, 2021). We bootstrap-sample the data 500 times, stratified over the make-model. We then assign treatment according to the method described in Section 5.2, and construct outcome variables conditional on cluster and individual assignments into treatment. We assume that direct and spillover treatment effects are linear and additive, except for the transaction outcome. [3] Spillovers are assumed to occur within the make-model cluster evenly for both treated and untreated posts. We assume no inter-cluster spillovers.

The outcome variables, which are standardized and identical to the primary outcomes described in Section 7, are the following:

- log(absolute difference between listing price and Price Calculator estimate)
- 1 if reported as sold
- log(self-reported transaction price)
- advertising index[4]
- buyer-attention index.

We estimate power of detecting the intend-to-treat (ITT) effects of direct treatment as well as cluster-level spillovers for a range of effect sizes (0.025 to 0.1 standard-deviation changes). We only measure pooled spillover effects (i.e. no saturation/intensity effect). We employ the same models as in the main analysis (i.e. logit for the binary outcome, and linear regressions for all other outcomes). These models include the same set of controls that we use for the main analysis (as discussed in Section 8.2). We report the false-discovery-rate-adjusted q-values based on 4 p-values corresponding to the main outcomes. These adjustments are made separately for direct treatment and spillover effects.

The results of power simulations are shown in Figure 4. We find that with controls on vehicle characteristics, we are able to detect treatment and spillover effects of 0.05 times the standard deviation for most primary outcome groups, both direct and spillover, with over 80% probability. The exception is the spillover effect of the binary transaction outcome, where we have sufficiently high power for an effect size just above 0.1 SD. These effect sizes translate roughly into 4,900 PKR (32 USD) in absolute difference between listing price and Price Calculator estimate (level mean: 138,000 PKR), 5 percentage-points in transaction

[3] Given that the transaction outcome is binary, we assumed that assignment into treatment would increase the probability of transaction by X%, where X is a standardized effect size based on the standard deviation of the binary variable.

[4] There is a minor difference in definitions of constituent variables, due to limitations in the pre-intervention data

probability (mean: 0.42), and 49,000 PKR (320 USD) in transaction price (level mean: 1,735,707 PKR).

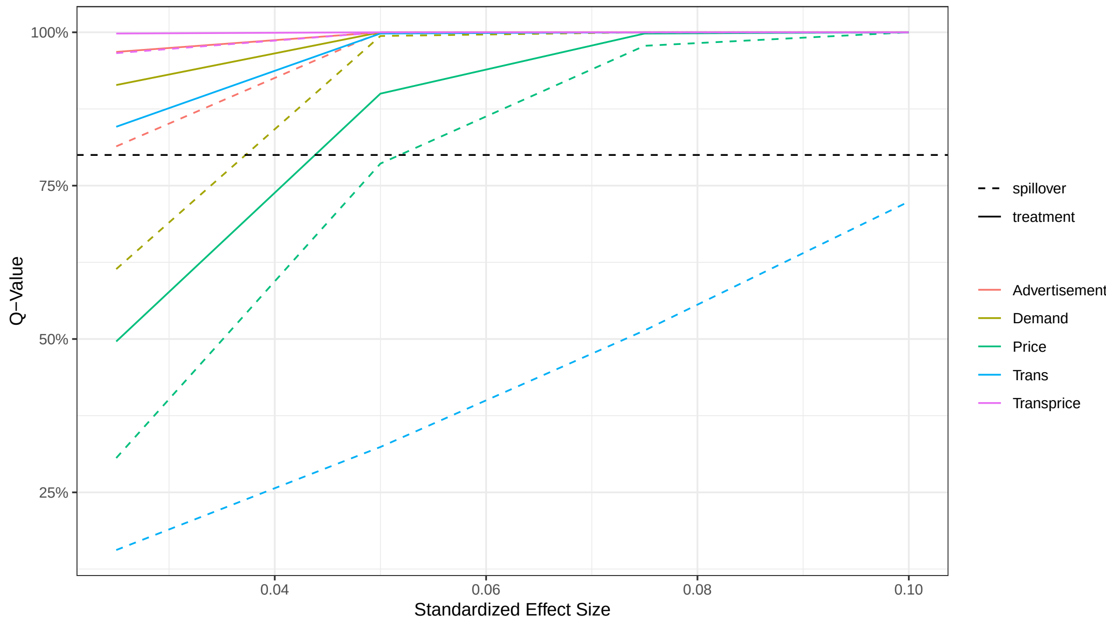

Figure 4: Power estimates of direct and spillover effects

6 Data

In this section, we highlight data items that we have used or intend to use for the design and main analysis. They are 1) vehicle characteristics and prices associated with the posts, 2) uses of advertising tools and other paid services, 3) views and clicks as a proxy for buyer interest, and 4) telephone survey data on transaction outcomes and their (unincentivized) beliefs of expected transaction prices. All data, except for the survey, come from PakWheel's near-live database. Other data items we use for design and secondary analysis include i) aggregated search engine results in terms of key words and their combinations, ii) daily search listing orders from PakWheels, and iii) a usage log of a previous iteration of the Price Calculator, which preceded the experiment.

6.1 Posts

PakWheels' database tracks every post on the platform. Once a post is created, it is vetted against spam or fraud, made publicly available on the platform, then removed after 90 days or once the user asks for it to be taken down. We collect the following measures from the database:

- timing of the post's creation, approval, and closure
- vehicle characteristics
  - basic information such as make, model, model-year, mileage, sellers' location, and registration city
  - additional information about vehicle characteristics, such as version, assembly, engine size and capacity
- listing price
- self-reported transaction outcome (e.g. sold to customer on the platform, sold through other means, decided not to sell)
- self-reported transaction price, if sold.

The database also tracks any updates to variables over the course of posts' active status. This allows us to capture sellers' initial choice of listing price before and after exposure to the Price Calculator estimate.

6.2 Advertising tools and vehicle inspection services

PakWheels' database also tracks users' platform credit purchases and usage, which we consider to be measures of sellers' advertising choices. Users on PakWheels have two primary tools for advertising: "bump" and "feature" credits. A "bump" credit allows sellers to bring their post to top of the result page in the default, reverse-chronological order. This effectively increases the post's visibility, as more people look at the first pages of listings. On the other hand, a "feature" credit would put their post in a few reserved spots at the top of the result page and labeled as a "featured ad", in a similar way as promoted ads on Google searches. Posts are otherwise listed in the reversed chronological order within the class of featured ads.

Another way for sellers to attract buyers' attention to their posts is to provide signals of vehicle quality. In order to do so, sellers can request in-person inspections by PakWheels' mechanics, who give scores (out of 100) on 8 dimensions (engine, brakes, suspension, interior, AC, electrical, exterior, and tires) based on a pre-specified rubric. The vehicles would pass the inspection if the unweighted average of scores over these 8 dimensions is above a threshold. They can then be marked as "PakWheels certified" on the platform for an additional fee. Given that certification is endogenous to vehicle quality, we use the data on whether or not the vehicle was ever inspected, as opposed to certified.

6.2.1 Expenditures on advertising tools

Another way of expressing sellers' advertising choices would be to in terms of amount paid to the platform for advertising. This is made difficult, however, by the fact that credits for bumps and features are purchased in bundles, and users can apply them to any posts that they own. We therefore do not plan on using this measure as a primary outcome. Nonetheless, we will collect data on advertising expenditures for robustness checks where the unit of analysis is the user. The data set contains information on purchase timing, descriptions of the bundles or services, quantity, and prices.

6.3 Buyer-attention measures

One of our hypotheses is that the price information intervention and resulting changes to pricing and advertising would affect buyers' attention to certain posts. In order to construct measures of buyer attention, we access PakWheels' data on views and clicks at the post level. We are able to collect cumulative measures of the following:

- page views (i.e. clicks on the post)
- clicks on "Show Phone Number" button within the post to contact the seller.

We also capture the number of times each post appears on search listings. We will run an analysis including this measure in the index as an alternative specification, as well.

6.4 Transaction survey

We use self-reported transaction prices to train and test the Price Calculator estimates. These self-reported data are collected in an online form whenever sellers choose take their posts down. There are concerns about the accuracy of reported transaction prices, particularly for the following reasons:

1. Transaction prices may be selectively reported (e.g. if those who fetched a higher prices are more likely to report).
2. Conditional on reporting, sellers may obfuscate the true value, leading to more noise in the price estimate.
3. Conditional on reporting, sellers may feel that the reported price should be inflated or deflated (because of, for example, their beliefs about a fair transaction price, or the desire to appear successful).
4. Conditional on reporting, sellers may find it easy just to repeat the listing price they have already given for their post.

5. Conditional on reporting, sellers may simply put a random number down to "get it out of the way."

We can identify the extent of number 4. by comparing transaction prices with listing prices, and address 5. via data cleaning. Concerns like number 2. may introduce noise but should not bias the Price Calculator estimate or our empirical analysis.

We are unable to directly address concerns 1. and 3. from the data, nor is it realistic for us to request sales receipts or access other independent sales records. Instead, we will select a subset of listings (stratified over the vehicle model) and conduct a short phone survey to the owners. We will survey up to 1,000 posts before the experiment as a pilot, then up to 3,000 posts during and after the intervention. The primary objective of the survey is to ask the following questions:

- if they have sold their vehicle
- the transaction price, if they have sold the vehicle
- price at which they expected to sell their vehicle, when they created the post.

Additionally, we will also ask the following questions for further secondary analysis we wish to conduct:

- reasons for not selling the vehicle (if not sold)
- relationship with the buyer (if sold)
- recollection of exposure to the Price Calculator, and of the estimates
- beliefs about the accuracy of the Price Calculator estimate
- what information/experience they relied on to set the list price
- difficulty of receiving enough inquiries, and/or good price offers.

PakWheels routinely conducts short customer telephone surveys for data quality assurance, and they have conducted some pilot surveys. The remainder, however, will be conducted by a Lahore-based economics research organization, the Institute of Development and Economic Alternatives (IDEAS).

6.5 Other data sources

Besides the main data items listed above, we collect additional types of data to a) validate our experimental design, b) conduct robustness checks, and c) run exploratory analysis on the structure of spillovers. The following is a list of such data items.

6.5.1 Search engine logs

Aggregated search engine logs tell us which combinations of terms are used most frequently by viewers on PakWheels. We use these aggregate statistics for our justifications for market cluster groupings. Our objective is to minimize concerns about inter-cluster interference, but also retain as many randomization clusters for the step as possible. Our aggregate search logs data are taken from the month of August, 2020. They represent tens of millions of searches over the month, and our data contain numbers of searches per combination of search terms (e.g. make, model, model-years in range, city, range of listing prices). We capture 35,000 most common search combinations, which account for 93% of all searches. We use these data for our definition of clusters in Section 5.2.1.

6.5.2 Listing orders

Beyond the primary analysis, in which we measure the average spillover effects on treated clusters, we plan to assess the extent to which the spillover effects depend on the "proximity" to treated posts, such as how close a given ad is to treated peers in ad listings. For this, we web-scrape listing orders in their default, reverse-chronological order on a daily basis for each make-model cluster in the sample.

6.5.3 Use of an old Price Calculator

We also track usage of a previous version of Price Calculator, which our intervention will replace. The previous iteration of the Price Calculator was designed and implemented prior to the beginning of our research collaboration with PakWheels. It was contained in a separate module in PakWheels' website and mobile apps, unintegrated with the posting process, and was discontinued at the end of December, 2020. The old Price Calculator offered predictions to only a handful of make-model-model years of certain colors, locations, and mileage. PakWheels keeps a log of all price estimates the old Price Calculator provided at each instance. This dataset contains user ID, search inputs (make, model, model year, location, mileage, if seller or buyer), the price estimates, and the time stamp.

7 Outcomes

7.1 Primary outcomes

Our primary outcomes for our analysis of direct treatment and spillover effects are the following:

- log(absolute difference between listing price and Price Calculator estimate)

- 1 if reported as sold
- log(self-reported transaction price)
- advertisement index
- buyer-attention index.

We define these outcomes in the subsections below.

7.1.1 Absolute difference in prices

We consider the effect on listing prices as the "first-stage" of our intervention, in that effects on other primary outcomes hinge on the changes to listing prices and their distributions. We also expect to observe adjustments in the listing price toward the optimal listing price, which would include some margin for expected bargaining. In order to capture this type of convergence, we define our primary price outcome to be the natural-log transformation of the absolute difference between the final listing price and the Price Calculator estimate. PakWheels calculate and provide the Price Calculator estimate only to treated posts; we will therefore estimate the prices that control posts *would have received*, based on the identical model specifications and parameters as the one used in the experiment.

As discussed in Section 6.1, sellers can update prices and other features of their posts throughout the course of their' active status. We therefore need to decide on at which point to record the listing price. We will use the listing price at the end of the post's active status, so that all the factors that may affect the listing price would be included. This is because direct effects of the Price Calculator estimate may happen when the post is created, while indirect effect may occur even after the post is created through feedback from buyers and competition with other posts. We will run robustness checks as well using the listing price from a) when the post is created and b) 7 days after the post's creation.

7.1.2 Transaction outcome and price

Sellers on PakWheels can take down their posts once they no longer wish to receive inquiries, or the post expires after 90 days since the initial posting. When the post is taken down, sellers are asked on PakWheels' online platform if they have sold their vehicles. They are required to respond in order to have their ads taken off. They are given options on the form (e.g. sold via PakWheels' website, sold via others, chose not to sell, etc.) and most sellers choose oe of them. However, some respond as "Other" yet report in the comment section that they have sold the vehicle. Our transaction outcome variable accounts for this to the best extent possible by string cleaning responses classified as "Other." The transaction variable is 1 if the seller reported a sale, 0 otherwise.

Sellers are also prompted to report the transaction price on the online form, if they report to have sold their vehicle. The value is missing for those who do not report their transaction outcome. We also remove inputs outside of reasonable price range for their given make-model. We use the natural log of transaction price as the outcome variable in analysis.

These self-reported outcome data are likely the best source of information on transactions and prices across a wide range of vehicle characteristics and locations in Pakistan. However, they may be vulnerable to biases and are checked against values collected via a telephone survey described in Section 6.4. We plan on using responses from this survey to construct analogous outcome variables for robustness checks.

7.1.3 Advertisement index

One of our primary hypotheses is that, when faced with novel price information, sellers adjust their strategic choices along two margins; list pricing and advertising. We will capture sellers' choices on advertising with data on paid services on PakWheels. As discussed in Section 6.2, sellers can increase visibility of their posts and/or signal quality by "bumping", "featuring," and requesting an inspection for their vehicle. In order to capture both intensive and extensive usage of advertising tools, we construct an index measure consisting of the following variables:

- number of "bumps" the seller applies to the post
- number of weeks the seller "features" the post
- 1 if the seller requests PakWheels to have the vehicle inspected.

7.1.4 Buyer-attention index

We also hypothesize that the price information intervention, and causal effects on pricing and advertising, affects treated posts' visibility on the platform. In order to capture this effect on post's visibility and buyer attention, we construct an indexed measure from data discussed in Section 6.3. The index consists of the following variables:

- page views (i.e. clicks on the post)
- clicks on "Show Phone Number" button within the post to contact the seller.

7.2 Secondary outcomes

7.2.1 Survey data

We classify the outcomes from our endline survey as follows:

1. validation of self-reported transaction outcome

- **1 if the vehicle is sold**
- **transaction price (if sold)**
- reasons for not selling the vehicle (if not sold)
- relationship with the buyer (if sold)

2. salience of the Price Calculator instrument
   - 1 if seller recalls seeing the Price Calculator estimate
3. mechanisms of acquiring price information and updating beliefs
   - **Expected transaction price at the time of initial posting**
   - 1 if seller seeks price information from an existing stock of listings
   - 1 if seller believes it is difficult to receive enough inquiries on PakWheels
   - 1 if seller believes it is difficult to receive acceptable price offers on PakWheels.

We consider the outcomes in bold as the main objectives of the survey. Rest of the outcomes have not been finalized as of June 13, 2021, as the pilot survey on these outcomes is yet to be completed.

7.2.2 Post's duration on the platform

PakWheels' database reports when each post is created and taken down, so we can calculate the duration of post's active status on the platform. One challenge is that posts may be left inactive for a period of time, so this would not be a measure of sellers' active participation in the market. This makes it difficult to interpret meaning of any causal effect on this variable other than in aggregate as a measure of market congestion. For this reason, this outcome is relegated as secondary.

7.2.3 Price changes, and convergence to estimated price

We choose the deviation of the listing price from the Price Calculator estimate as a primary outcome variable. It is possible, however, that the treatment effect on the choice list price may be better captured if the Price Calculator induces a level shift in price or an extensive-margin change on whether or not sellers ever adjust their listing prices. It is also possible that the treatment effect on the listing price is asymmetrical around the Price Calculator estimate. We intend to address these possibility with following alternative outcomes pertaining to the listing price:

- log(list price)
- 1 if the listing price is ever modified
- difference between the initial and final listing prices.

7.2.4 Cluster-level outcomes

Our two-stage randomization procedure allows us estimate impact on cluster-level outcomes, since treatment is phased in over groups of clusters over time. Part of our secondary analysis focuses on cluster-level aggregate measures of moments of prices, page views, and post duration. We construct the following variables at the cluster-week level:

- number of new posts
- number of active posts
- standard deviation and kurtosis of the listing price
- standard deviation and kurtosis of page views.

7.2.5 Spillovers based on listing order

As discussed in Section 6.5.2, we web-scrape the listing order from PakWheels. We will use these data to a) estimate the impact of treatment on sellers' positioning on the search listings, and b) create a measure of potential spillovers' intensity based on proximity to treated posts.

We define outcomes from PakWheels' scraped data as follows:

- number of days a post is on the first page of the make-model level search result
- average page number of the search results over the course of its active status.

We also construct the following variables as a proxy of exposure to treated posts:

- number of days adjacent posts are treated with a Price Calculator estimate
- average number of treated posts within the listing result page, over the course of the post's active status.

8 Empirical analysis

8.1 Balancing checks

We will not be able to observe baseline outcomes of the experimental sample, because we will be treating a *flow* of new posts and our outcomes are what happens to those posts. Given this limitation, we test balance over our outcome variables using listings data from a pre-treatment period. We will apply the same sample selection criteria and randomization procedure as the experiment, described in Section 5.2. This ensures that the we have balance over our primary outcome variables on pre-treatment-period data. We can check further balance on a) primary outcomes that were not included in the randomization procedure due to concerns about data quality or availability, b) constituent variables to outcome indices, c) secondary outcome variables, and d) other user-level characteristics.

8.2 Primary analysis

8.2.1 Outcomes of interest

In our primary analysis, we will estimate intend-to-treat and treatment-on-treatment effects of both direct treatment and spillovers. The following are our primary outcomes and corresponding research questions highlighted in Section 2:

- absolute difference between listing price and Price Calculator estimate (Research question 1.1.)
- binary transaction outcome (Research question 1.2.)
- log(transaction price) (Research question 1.2.)
- indexed measure of advertisement usage (Research question 1.4.)
- buyer-attention index (Research question 2.3.)
- spillover effect on all outcomes above (Research question 2.1.).

8.2.2 Main specification: Intend-to-treat

We estimate the intent-to-treatment effect of being provided the Price Calculator estimate, and of being in the spillover make-model clusters. Coefficients of interest are those of assignment and cluster variables, as the following estimating equation shows:

$$\begin{aligned} Y_{i,p} = \beta_0 + \beta_1 * Assign_{i,m,w} + \beta_2 * Cluster_{m,w} \\ + \psi_w + \gamma_m + X_p^{'} \rho + \omega_m + \epsilon_{i,p} \end{aligned} \tag{1}$$

The subscripts used in the equation above indicate the following:

- i: individual user identifier (defined by PakWheel's user ID)
- p: post (multiple posts could belong to a given i)
- m: vehicle make-model cluster

- w: posting week.

$\hat{\beta}_1$ and $\hat{\beta}_2$ capture the ITT effects. *Assign* is the binary direct treatment variable, and *Cluster* is the cluster-week level dummy variable, which equals 1 if that cluster-week is selected for assignment in the first stage of randomization. Our hypotheses, based on the theoretical framework in Section 4, are that the treatment would:

- reduce the absolute difference between listing price and Price Calculator estimate for directly treated posts ($\beta_1 < 0$) as well as spillover posts ($\beta_2 < 0$)
- increase the transaction probability for directly treated posts ($\beta_1 > 0$), and increase of decrease it for spillover posts ($\beta_2 \neq 0$)
- increase or decrease the transaction price for directly treated posts ($\beta_1 \neq 0$) as well as spillover posts ($\beta_2 \neq 0$), conditional on the occurrence of transaction
- weakly increase the indexed measure of advertisement usage for directly treated posts ($\beta_1 \geq 0$) as well as spillover posts ($\beta_2 \geq 0$)
- increase the buyer-attention index for directly treated posts ($\beta_1 > 0$), and increase or decrease it for spillover posts ($\beta_2 \neq 0$).

For all dependent variable other than the binary transaction outcome, we will use linear regressions. For the binary outcome variable, we will use the logit model. We will cluster the error at the make-model level (with the error term ω_m shown in Equation 1 representing this assumption), as the first stage of the randomization is conducted at this level. We will also estimate these models using robust standard errors to in order to make sure that statistical significance do not change meaningfully.

8.2.3 Main specification: Treatment-on-the-treated

As discussed in Section 5.3, we expect some level of treatment non-compliance based on the fact that Price Calculator estimates cannot be delivered to apps that are not updated during the intervention period. This non-compliance is likely non-random, so we plan to instrument for treatment take-up using the assignment variable.

The treatment-on-the-treated (TOT) effect is estimated via 2SLS, with *Assign* instrumenting for *Treat*, and *Cluster* instrumenting for itself.

$$\begin{aligned} Y_{i,p} = \theta_0 + \theta_1 * \widehat{Treat}_{i,p} + \theta_2 * Cluster_{m,w} \\ + \psi_w + \gamma_m + X_p^{'}\rho + \omega_m + \epsilon_{i,p} \end{aligned} \tag{2}$$

The first-stage specification for $\widehat{Treat}$ is as follows:

$$Treat_{i,p} = \phi_0 + \phi_1 * Assign_{i,m,w} + \psi_w + \gamma_m + X_p^{'}\tau + \omega_m + \epsilon_{i,p} \quad (3)$$

$\hat{\theta_1}$ and $\hat{\theta_2}$ capture the TOT effects. Other than that we instrument for treatment with assignment, our identification strategy is identical to that of the ITT effects.

An alternative way of identifying the spillover effect may be to estimate the relationship between spillover outcomes and the *number* of treated posts in their cluster. This would make sense if we were to exogenously vary the intensity of treatment varies across models, but that is not our design. Instead we will focus on the average spillover effect to treated market cluster, and then explore the structure of spillovers as secondary analysis, as shown in Section 8.3.6.

8.3 Secondary Analysis

We exclude all tests in the secondary analysis from corrections for multiple hypothesis testing, in order to conserve power for the primary analysis.

8.3.1 Level shifts and changes to the listing price (Research Questions 1.1.1. and 1.1.2.)

As discussed in Sections 7.1.1 and 8.2, we use the deviation of listing prices from Price Calculator recommendations as the primary price outcome. However, we are also interested in any level shifts, or any changes at all, that the Price Calculator intervention might induce on the listing price. We therefore plan on estimating both ITT and TOT effects on the following outcomes using Equation 1:

- log(final listing price) (Research Question 1.1.1.)
- 1 if listing price was ever updated (Research Question 1.1.2.)
- absolute difference between the initial and final listing prices.

8.3.2 Heterogeneous treatment effect, and mechanisms of spillovers (Research Questions 1.1.3. and 2.2.)

Our theoretical framework (see Section 4.1) indicates that spillovers may occur via multiple channels; sellers may adjust their choices based on changes in their beliefs about market demand (as information itself spills over), the distribution of page views, or the match rate between buyers and sellers (theoretical prediction 4.). Similarly, the magnitude of direct treatment may vary depending on the salience of information and the extent of search frictions and congestion. These predictions on mechanisms are framed as Research Questions 1.1.3. and 2.2. in Section 2. We will test these hypotheses by identifying how the sizes of direct treatment effect and spillovers differ by focusing on the following four dimensions of heterogeneity:

1. price information prior to intervention: dummy variable for the previous availability of an old Price Calculator (discontinued in December, 2020)
2. sellers' experience : indexed measure of seller characteristics, consisting of:
   - number of previous posts
   - months since first post on PakWheels
3. information frictions: pre-intervention variation of listing prices at the make-model level, conditional on Price Calculator estimates
4. search frictions and congestion: indexed measure at the make-model cluster, consisting of:
   - number of posts per month
   - average duration of posts' active status
5. Product heterogeneity: make-model-level index that captures variety in the following characteristics:
   - model years
   - versions
   - color

We modify Equation 1 to get the following for the same set of outcomes as those listed in Section 8.2. We interact both the direct treatment and spillover terms with these measures of heterogeneity, denoted as $H_{i,p}$.

$$\begin{aligned} Y_{i,p} = \beta_0 + \beta_1 * Assign_{i,m,w} + \beta_2 * Cluster_{m,w} \\ + H_{i,p} * \beta_3 * Assign_{i,m,w} + H_{i,p} * \beta_4 * Cluster_{m,w} \\ + H_{i,p} * + \psi_w + \gamma_m + X_p^{'} \rho + \omega_m + \epsilon_{i,p} \end{aligned} \tag{4}$$

$H_{i,p}$ varies at the following levels:

1. pre-intervention availability of price information: individual post level, depending on vehicle characteristics (vehicle make, model, model-year, color, and city)

2. sellers' experience : seller level
3. information frictions: make-model level
4. search frictions and congestion: make-model level.

For variables that only vary at the make-model level, we do not include the uninteracted $H_{i,p}$ term as it would be collinear with the market fixed effect γ_m.

8.3.3 Market power

Existing body of evidence suggests (e.g. Mitra et al. (2018)) that the pass-through of price information may be dependent on the extent of market power. We therefore plan on testing how market power drives the effect of price information. We will identify heterogeneous treatment effects based on the following variables

- the Herfindahl–Hirschman Index (HHI) of the posts' ownership at the make-model level
- the share of professional dealers at the make-model level.

8.3.4 Changes in beliefs (Research Question 1.3.)

As shown in Section 4.1, what we believe to be a key mechanism of treatment is through changes in sellers' beliefs about the demand distribution; tailored price information leads sellers to have less noisy beliefs about the eventual transaction price. In order to test this mechanism, we estimate the treatment effect on sellers' non-incentivized beliefs on the vehicle's expected transaction price, which we will collect through a phone survey (see 6.4). The specifications are identical to that of Section 8.2.

8.3.5 Impact on congestion (Research Question 2.4.)

Another potential mechanism of spillovers and other knock-on effect is through changes in the number of active posts in a market segment, affecting the extent of congestion and search friction. The price information intervention may improve the matching process if, for instance, sellers and buyers are able to find better matches through less noisy listing prices. This may in turn affect the extent of congestion in treated market segments. Though not formally incorporated (yet) into our theoretical framework, we intend to test this hypothesis by estimating the treatment and spillover effects on the duration of active status on the platform, at the post level. The econometric specification is identical to one listed in Section 8.2.

8.3.6 Alternative specifications of spillover effects

In our default specifications, we measure spillover effects as an average over cluster-weeks that are assigned into treatment in the first step of the randomization procedure. It is possible, however, that spillover effects are more localized in the ad-listing space around treated posts, or that there needs to be a high concentration of treated posts in the ad-listing space for there to be spillovers to untreated posts. In order to capture effects of these alternative functional forms, we construct variables for potential spillover intensity, using the daily listing order data. We will construct the following variables at the individual post level:

- Share of other posts within the same result page that received the Price Calculator estimates (averaged, from daily ad-listing orders)
- Number of days in the first week of the post's active status in which posts right above or below are treated.

We will estimate the following equation, where the alternative spillover variables go into $S_{i,p}$

$$\begin{aligned} Y_{i,p} = \beta_0 + \beta_1 * Assign_{i,m,w} + \beta_2 * S_{i,p} + \beta_3 * S_{i,p} * Assign_{i,m,w} \\ + \psi_w + \gamma_m + X_p^{'}\rho + \omega_m + \epsilon_{i,p} \end{aligned} \tag{5}$$

β_1 denotes the direct treatment effect, β_2 the spillover effect based on the intensity measures above, and β_3 additional spillovers on the treated. We no longer include the cluster-level assignment variable, but will control for the cluster fixed effect.

8.4 p-value adjustments

In order to address the issue of multiple hypothesis testing, we follow Romano and Wolf 2005 and correct for the false discovery related to tests on the following primary outcomes:

- log(absolute difference between listing price and Price Calculator estimate)
- binary transaction outcome
- log(transaction price)
- indexed measure of advertisement usage
- buyer-attention index.

Given that we consider direct treatment and spillovers as separate hypotheses, we should adjust their critical values separately.We will therefore report adjusted critical values based on the Romano-Wolf procedure for 2 sets of 5 tests. We will then report unadjusted p-values for secondary outcomes.

References

Jenny Aker. Information from Markets Near and Far : Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics*, 2(3):46–59, 2010.

Jenny Aker and Isaac Mbiti. Mobile Phones and Economic Development in Africa. *Journal of Economic Perspectives*, 24(3):207–232, 2010. ISSN 00220388. doi: 10.1080/00220388.2012.709615.

Treb Allen. Information Frictions in Trade. *Econometrica*, 82(6):2041–2083, 2014. doi: 10.3982/ecta10984.

Tahir Andrabi, Jishnu Das, and Asim Ijaz Khwaja. Report cards: The impact of providing school and child test scores on educational markets. *American Economic Review*, 107(6): 1535–1563, 2017. ISSN 00028282. doi: 10.1257/aer.20140774.

David Atkin and Dave Donaldson. Who's getting globalized? The size and implications of intra-national trade costs. 2015.

Jie Bai, Maggie Xiaoyang Chen, Jin Liu, and Daniel Yi Xu. Search and Information Frictions on Global E-Commerce Platforms: Evidence from AliExpress. 2020.

Michael R. Baye, John Morgan, and Patrick Scholten. Information, Search, and Price Dispersion. In T. Hendershott, editor, *Handbook on Economics and Information Systems*, pages 323–375. Elsevier, 2007. doi: 10.1016/s1574-0145(06)01006-3.

Christopher Blattman and Laura Ralston. Generating Employment in Poor and Fragile States: Evidence from Labor Market and Entrepreneurship Programs. 2015.

Victor Couture, Benjamin Faber, Yizhen Gu, and Lizhi Liu. E-Commerce Integration and Economic Development: Evidence from China. 2018.

Peter A Diamond. Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy*, 90(5):881–894, 1982.

Michael Dinerstein, Liran Einav, Jonathan Levin, and Neel Sundaresan. Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7): 1820–1859, 2018. ISSN 00028282. doi: 10.1257/aer.20171218.

Liran Einav, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7(2):215–247, 2015. ISSN 19457685. doi: 10.1257/mic.20130046.

A. Nilesh Fernando, Gabriel Tourek, and Niharika Singh. Hiring Frictions in Small Firms: Evidence from an Internet Platform-based Experiment. 2020.

Andrey Fradkin. Search Frictions and the Design of Online Marketplaces. 2015. doi: 10.4108/eai.8-8-2015.2260850.

Ali Hasanain, Muhammad Yasir, and Khan Arman. No bulls : Experimental evidence on the impact of veterinarian ratings in Pakistan. 2019.

John J. Horton. Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. *Management Science*, 65(8):3518–3540, 2019. ISSN 15265501. doi: 10.1287/mnsc.2018.3116.

Robert Jensen. The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *The Quarterly Journal of Economics*, 122 (3):811–846, 2007. doi: 10.1080/02724980343000242.

Robert Jensen and Nolan H. Miller. Market integration, demand, and the growth of firms: Evidence from a natural experiment in India. *American Economic Review*, 108(12):3583–3625, 2018. ISSN 19447981. doi: 10.1257/aer.20161965.

Dahyeon D J Jeong. Creating (Digital) Labor Markets in Rural Tanzania. 2020.

David McKenzie and Christopher Woodruff. What are we learning from business training and entrepreneurship evaluations around the developing world? *World Bank Research Observer*, 29(1):48–82, 2014. ISSN 02573032. doi: 10.1093/wbro/lkt007.

Sandip Mitra, Dilip Mookherjee, Maximo Torero, and Sujata Visaria. Asymmetric information and middleman margins: An experiment with Indian potato farmers. *Review of Economics and Statistics*, 100(1):1–13, 2018. ISSN 15309142. doi: 10.1162/REST_a_00699.

Ryan T. Moore. Multivariate continuous blocking to improve political science experiments. *Political Analysis*, 20(4):460–479, 2012. ISSN 10471987. doi: 10.1093/pan/mps025.

Joseph P. Romano and Michael Wolf. Stepwise multiple testing as formalized data snooping. *Econometrica*, 73(4):1237–1282, 2005. ISSN 00129682. doi: 10.1111/j.1468-0262.2005.00615.x.

Meredith Startz. The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade. 2018. doi: 10.2139/ssrn.3096685.

George J. Stigler. The Economics of Information. *Journal of Political Economy*, 69(9):213–225, 1961. ISSN 1098-6596.

A Theoretical Framework

A.1 Introduction

We present a simple conceptual framework to demonstrate how sellers facing information and search frictions set listing prices, promote their posts by advertising, and respond to information about market conditions. The model is a static search framework, deriving inspiration from established work such as Stigler (1961) and Diamond (1982). Most canonical models that focus on the effect of access to price information and assume full knowledge of parameters on market friction, demand distributions, etc. (Baye et al. 2007). We, however, introduce the following deviations from a standard search models:

- Sellers have biased or noisy beliefs about the distribution of buyer willingness-to-pay (WTP).
- This, along with noisy beliefs about the match rate and efficacy of advertising, may lead to biased or noisy beliefs about the probability of sale and to suboptimal pricing.
- Sellers can influence the match rate with potential buyers by engaging in costly actions, i.e. advertising.

We first introduce the model without information friction, and consequently with noiseless beliefs. The remainder of this section is organized as follows; Section A.2 lays out the set-up of our model and provides definitions on terms and parameters. Section A.3 defines the objective function and the maximization problem in terms of the listing price and advertising choices. Section A.4 gives optimality conditions in the case of no information friction. Section A.5 shows how individual choices may be altered when there is noise in beliefs about the demand. Section A.6 concludes by providing predictions on the role of information frictions and noisy beliefs on demand, in terms of sellers' choice variables and transaction outcomes.

A.2 Set-up

Suppose that we have a seller i, who is endowed with asset- and seller-characteristics s_i and information set $\mathcal{I}_i$.

- Seller i forms a belief about the demand distribution for their asset based on the information set $\mathcal{I}_i$.
- Some sellers are provided with the Price Calculator estimate denoted as $\mathcal{T}_i$, which improves the accuracy of information set $(\mathcal{I}_i \cup \mathcal{T}_i)$.
- Treated sellers update their beliefs about demand distribution, based on the new information set.

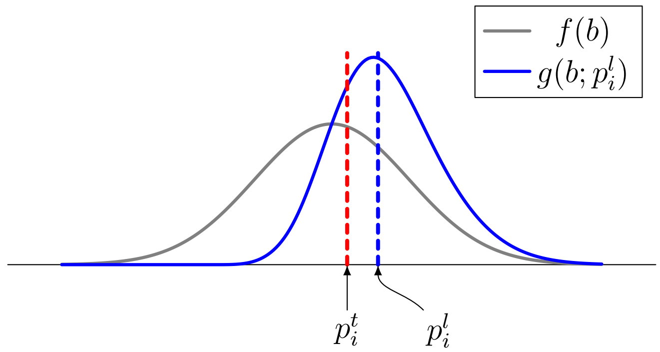

Figure 5: Relationship between list price, the buyer it draws, and the TIOLI price

- Seller i chooses a listing price p_i^l and amount of advertisements a, based on s_i and $\mathcal{I}_i$.
- Choices of p_i^l and a affect the distribution of potential buyers with whom seller i is matched via a Poisson process.
- Once a match occurs, seller i makes a take-it-or-leave-it (TIOLI) offer p_i^t below p_i^l to the potential buyer.
- Transaction occurs if matched buyer's WTP is higher than p_i^t.

We denote the CDF of true buyer WTP as $F(b)$, and the distribution of potential buyers that get matched to the seller, conditional on p_i^l, as $G(b; p_i^l)$. The distinction between $F(b)$ and $G(b; p_i^l)$ is key, since we assume that the seller's choice of the listing price p_i^l skews the distribution of potential buyers (who may click on the post depending on the listing price) towards itself, p_i^l. The relationship between p_i^l and $G(b; p_i^l)$ can be described schematically in Figure 5:

A.3 The objective function

Under no information friction, seller i chooses the listing price and advertisements to maximize the following:

$$V(p_i^l, a; s_i) = -c - k(a) + \gamma(a) \int_{b \in G} \max_{p_i^t} [\mathbb{E}\pi(p_i^t; p_i^l, s_i)] dG(b; p_i^l) \tag{6}$$

Sellers incur a constant cost of search, denoted as c. They also incur a variable cost $k()$, based on the amount spent on advertising, a. The term $\gamma(a)$ is a Poisson match rate between a seller and a potential buyer, and it is an increasing function with respect to a. We denote the seller's utility from transaction as $\pi(p_i^t, s_i)$. This is not strictly a profit term, as sellers may also have preferences over how quickly to sell the vehicle (i.e. a function of s_i).

A.3.1 Setting the TIOLI price p_i^t

Note that there is a distinction between the listing price p_i^l and the (TIOLI) offer price p_i^t. We assume that there is one-to-one correspondence between them conditional on seller characteristics. We further assume that potential buyers cannot perfectly infer p_i^t from p_i^l, because this depends on the seller's individual characteristics s_i as well as $\mathcal{I}_i$. This allows us to simplify the seller's problem as maximization of $V(p_i^l, a; s_i)$ via choices of p_i^l and a. We get the following first-order condition over the choice variable p_i^l:

$$\pi'(p_i^t(p_i^l, s_i))[1 - G(p_i^t; p_i^l)] - g(p_i^t; p_i^l)\pi(p_i^t(p_i^l, s_i)) = 0 \tag{7}$$

The objective function can be simplified to the following:

$$V(p_i^l, a; s_i) = -c - k(a) + \gamma(a)\pi(p_i^t(p_i^l, s_i))[1 - G(p_i^t(p_i^l, s_i))] \tag{8}$$

A.4 Identifying optimal p_i^l and a

Taking the first-order condition of Equation 8 with respect to p_i^l gives the following expression, where we see that the choice of optimal p_i^l is independent of a under no information friction.

$$0 = \frac{dV}{dp_i^l} = \gamma(a)[\pi'(p_i^t)\frac{dp_i^t}{dp_i^l}[1 - G(p_i^t(p_i^l, s_i))] - \pi(p_i^t(p_i^l, s_i))\frac{dG(p_i^t)}{dp_i^t}\frac{dp_i^t}{dp_i^l}] \tag{9}$$

Rearranging Equation 9, we get:

$$[1 - G(p_i^t(p_i^l, s_i))]\pi'(p_i^t)\frac{dp_i^t}{dp_i^l} = \frac{dG}{dp_i^t}\frac{dp_i^t}{dp_i^l}\pi(p_i^t(p_i^l, s_i)) \tag{10}$$

The left-hand side of Equation 10 is an expression of "marginal benefit" of price adjustment, i.e. the marginal change in the seller's payoff ($\pi'(p_i^t)\frac{dp_i^t}{dp_i^l}$) times the probability that a matched buyer accepts the TIOLI price ($1 - G(p_i^l, s_i)$). The right-hand side is an expression of the "marginal cost" of price adjustment, i.e. the marginal effect of the changes in listing price on the probability of TIOLI price's acceptance ($\frac{dG}{dp_i^t}\frac{dp_i^t}{dp_i^l}$) times the payoff ($\pi(p_i^t(p_i^l, s_i))$).

Similarly, taking the first-order condition of Equation 8 with respect to a gives the following expression that identifies the optimal a is conditional on a choice of p_i^l.

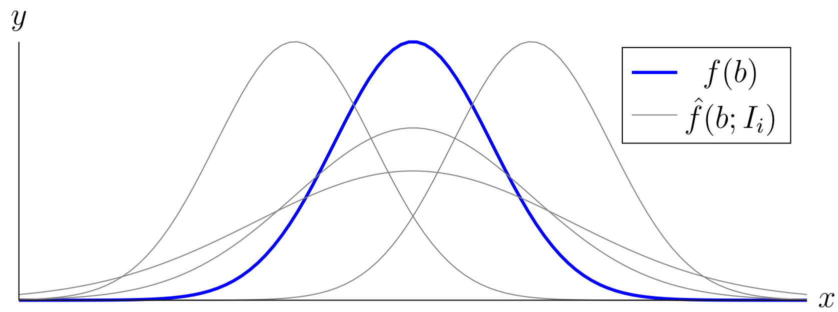

Figure 6: Noise in beliefs about the demand distribution

$$0 = \frac{dV}{da} = -k'(a) + \frac{d\gamma}{da}\pi(p_i^t(p_i^l, s_i))[1 - G(p_i^l, s_i)] \tag{11}$$

There is unique solution that satisfies Equation 11 if the Poisson match rate function $\gamma()$ is concave, and the cost function $k()$ convex. A component of the right hand side of Equation 11 is the marginal gain from advertising, which a product of changes in the Poisson match rate ($\frac{d\gamma}{da}$) and expected payoff ($\pi(p_i^t(p_i^l, s_i))[1 - G(p_i^l, s_i)$). This marginal gain is equal to the marginal cost of advertising, i.e. $k'(a)$.

A.5 Information friction and beliefs

The optimality conditions above hinge on the assumption that sellers have accurate beliefs about buyers' WTP and the Poisson match rate. However, if there is noise in sellers' beliefs about buyers' WTP (as shown in Figure 6) and the match rate, how would variation in information pertaining to WTP affect sellers' decisions?

Suppose that a seller holds a belief over $F()$, $G()$, and $\gamma()$. We denote one's beliefs as $\hat{F}(;\mathcal{I}_i)$, $\hat{G}(;\mathcal{I}_i)$, and $\hat{\gamma}(;\mathcal{I}_i)$, where $\mathcal{I}_i$ denotes information quality. The resulting optimality conditions then simply replace F with $\hat{F}(;\mathcal{I}_i)$, G with $\hat{G}(;\mathcal{I}_i)$, and γ with $\hat{\gamma}(;\mathcal{I}_i)$.

The idea behind our intervention is that by randomized provision of Price Calculator estimates, we exogenously improve the quality of information from $\mathcal{I}_i$ to $\mathcal{I}_i \cup \mathcal{T}_i$ such that we reduce deviations of $\hat{F}(;I_i)$, $\hat{G}(;I_i)$, and $\hat{\gamma}(;I_i)$ from their equivalents under no information friction.

A.6 Model predictions

A.6.1 Information intervention reduces variance in p_i^l

The optimality condition for p_i^l in Equation 10, under noisy beliefs, can be rearranged and expressed as follows.

$$\frac{\pi'(p_i^t)}{\pi(p_i^t(p_i^l, s_i))} = \frac{\frac{d\hat{G}}{dp_i^t}}{[1-\hat{G}(p_i^t(p_i^l(s_i, \mathcal{I}_i)))]} \tag{12}$$

Following the logic from Equation 10, Equation 12 shows that the seller sets their listing price p_i^l such that their *beliefs* about the expected payoff equals their *beliefs* about the cost. Their choice of p_i^l given information friction, however, does not necessarily satisfy the optimality condition from Equation 10. In other words, it is generally true that given a choice of p_i^l made under information friction (with access only to $\mathcal{I}_i$):

$$\frac{\frac{d\hat{G}}{dp_i^t}}{[1-\hat{G}(p_i^t(p_i^l(s_i, \mathcal{I}_i)))]} \neq \frac{\frac{dG}{dp_i^t}}{[1-G(p_i^t(p_i^l(s_i, \mathcal{I}_i)))]} \tag{13}$$

We do not make further assumptions about the structure of information friction. Rather, our point is that the choice of p_i^l made under noisy beliefs about the demand differs from one chosen under *no* noise in beliefs. In other words, if the Price Calculator intervention exogenously reduces noise in beliefs such that $\hat{G}(b; I_i)$ approximates $G(b; \mathcal{I}_i)$, then it would reduce variance in the distribution of p_i^l, as well. The conclusion from this section is summarized as a prediction below:

- Prediction 1.: Information intervention brings p_i^l closer to what it would be under no noise in beliefs about demand. (Research question 1.1.)

A.6.2 Information intervention increases expected payoffs

Although we do not make assumptions about the structure of information friction and resulting noise in beliefs, it is apparent that the expected payoff from participating in the market increases as information friction and noise in beliefs are reduced. As shown in Section A.4, the choice of p_i^l is made independently of that of a (but not vice versa). Given that the information intervention reduces noise and brings $\hat{G}$ closer to G, we can show that:

$$\pi(p_i^t(p_i^l, s_i))[1 - G(p_i^t(p_i^l(si, \mathcal{I}_i \cup \mathcal{T}_i)))] \geq \pi(p_i^t(p_i^l, s_i))[1 - G(p_i^t(p_i^l(si, \mathcal{I}_i)))] \tag{14}$$

This leads to the next prediction of our theoretical framework, that:

- Prediction 2.: Information intervention increases expected returns from the platform. (Research question 1.2.)

A.6.3 Information intervention reduces variance in a

We have so far shown that the choice of listing price can be affected by noise in sellers' beliefs about the demand, and that an information intervention would improve their payoff from engaging with the marketplace. How would their choice of advertising, then be affected by information friction? From Equation 11, we see that under no information friction, sellers use advertising up to the point where the expected marginal benefit of its use equals its marginal cost. Under information friction, however, sellers consume advertising tools to the point where their *beliefs* about the expected marginal benefit equals marginal cost. The following equation makes this point by modifying Equation 11, and putting the term corresponding to beliefs about expected payoffs in a wide hat:

$$\frac{d\gamma}{da(s_i, \mathcal{I}_i)} \widehat{\pi(p_i^t(p_i^l(s_i, \mathcal{I}_i)))[1 - G(p_i^t(p_i^l(s_i, \mathcal{I}_i)))]} = k'(a(s_i, \mathcal{I}_i)) \tag{15}$$

We, again, do not make assumptions about the structure of beliefs nor how they are updated. However, if the information intervention improves sellers' expected payoffs (Equation 14) *and* that sellers themselves believe this result enough to have their *beliefs* about the expected payoffs updated, then we should see that:

$$\frac{d\gamma}{da(s_i, \mathcal{I}_i)} \widehat{\pi(p_i^t(p_i^l(s_i, \mathcal{I}_i \cup \mathcal{T}_i)))[1 - G(p_i^t(p_i^l(s_i, \mathcal{I}_i \cup \mathcal{T}_i)))]} \geq k'(a(s_i, \mathcal{I}_i)) \tag{16}$$

Then, it follows that $a(s_i, \mathcal{I}_i \cup \mathcal{T}_i) \geq a(s_i, \mathcal{I}_i)$. In other words:

- Prediction 3.: Information intervention increases a if sellers' beliefs about expected returns from the platform are adjusted upward.(Research question 1.4.)

A.6.4 Spillovers and their mechanisms

The optimality conditions and predictions above are based on the assumption that exogenous shocks in a form of an information intervention affects only individual choices. What we speculate, in terms of spillover effects, is that individual's access to information and their choices may generate the following effects:

- Treated sellers' choices of p_i^l may generate changes to the quality of information signals available in the market, therefore affecting $\mathcal{I}_i$, where $i \neq j$.
- It is possible that buyer attention in market segments are zero-sum. If treated sellers' choices of p_i^l and a make their posts more appealing to potential buyers, then it may affect the distribution of buyers to sellers conditional on their list prices, i.e. $G()$.

- If the treatment affects the distribution of buyer attention, or the speed at which sellers and buyers are matched, then $\gamma()$ would be altered.

In other words, our predictions of the spillovers and their mechanisms are the following:

- Prediction 4.: Spillover effects could occur through reduction in noise in beliefs about demand, changes in the distribution of pageviews within treated market clusters, or improvement in the Poisson match rate.(Research questions 2.1., 2.2., and 2.4.)