The Limits of Centralized Pricing in Online Marketplaces and the Value of User Control*

Apostolos Filippas
Fordham

Srikanth Jagabathula
NYU

Arun Sundararajan
NYU

February 1, 2022

Abstract

We report experimental and quasi-experimental evidence from a "sharing economy" marketplace that transitioned from decentralized to centralized pricing. Centralized pricing increased the utilization of providers' assets, resulting in higher revenues but also higher transaction costs. Providers who were barred from accessing the price system made non-price adjustments, including reducing the availability of their assets, canceling booked transactions, and exiting the market. Providers who retained partial pricing control reacted substantially less but experienced similar revenue increases. We highlight the challenges of implementing centralized pricing and assessing its welfare effects. We show that partial control can mitigate these challenges, allowing providers to express their private and heterogeneous preferences while maintaining the benefits of centralization.

JEL Codes: L11, D47

*The authors thank the technology and management teams of the platform for sharing anonymized data. We thank four anonymous referees for their valuable comments and suggestions. Author contact information and code are currently or will be available at `http://apostolos-filippas.com`.

1 Introduction

A fundamental question in the design of online platforms is how to best set prices. Early online marketplaces, such as eBay and Amazon, delegated price-setting to users and implemented a variety of decentralized mechanisms including posted prices, haggling, and auctions. Today, technological advances have enabled some online platforms to take on price-setting: for example, peer-to-peer lending platform Prosper determines interest rates for loans, and ride-hailing platforms Uber and Lyft set prices for rides centrally and in real time.[1]

The rationale for centralized pricing is that platforms possess superior market supply and demand information than sellers. From an efficiency standpoint, centralized pricing promises to be particularly beneficial if market participants have homogeneous preferences, and if the goods and services traded are highly "perishable" (Cramer and Krueger, 2016). Furthermore, centrally set prices may offer users a more uniform and seamless experience, allowing platforms to better compete against traditional firms. However, absent a decentralized price system, platforms face the well-known challenges of eliciting private information, which can be particularly important when users' preferences are heterogeneous (Hayek, 1945). Barred from accessing the price system, users may also make non-price adjustments with potentially deleterious consequences (Hirschman, 1970; Hall, Horton and Knoepfle, 2018). Determining whether platforms may benefit from centralizing pricing requires a firm understanding of the challenges involved in doing so. This paper sets to identify some of these challenges, to elucidate their root causes, and to propose and evaluate a potential solution.

We report the results of an experiment that took place on a "sharing economy" marketplace for vehicle rentals during a transition from decentralized to centralized price-setting. Providers (owners/sellers) were randomly assigned to one of two versions of centralized pricing. In the first treatment group (T1), the platform assumed complete control over pricing the providers' vehicles. In the second treatment group (T2), the platform assumed partial pricing control: while prices were determined by the platform, providers could raise or lower the centrally set prices by up to 30 percent. Providers in a control group (T0) remained in the status-quo pricing, maintaining complete control over price-setting.

Providers in the treatment groups reacted to the pricing change by exiting the platform, reducing the market availability of their assets, and canceling booked transactions. Compared to the control group, providers in treatment group T1 exited at a 30.5% higher rate, reduced the availability of their vehicles by 18.9%, and canceled 19.8% more transactions. Providers in treatment group T2 reacted similarly, but the magnitude of their responses was significantly

[1] These technological advances include the wide-spread adoption of internet-enabled mobile devices, the decreases in the cost of creating, storing, and processing data, improvements in the design of algorithms, and the proliferation of business experimentation (Varian, 2010). In addition, market designers can draw upon more than 20 years of accumulated industrial experience in building online platforms and solving their fundamental problems (Filippas, Horton and Zeckhauser, 2020).

smaller: they exited at a 50.7% lower rate, made their assets 54.7% more available, and canceled 36.4% fewer transactions than providers in T1.

We next examine possible reasons for the responses of the providers. We find that centralized pricing increased providers' revenues, but lowered prices and increased the utilization of their vehicles. Providers in T1 experienced 53.2% higher revenues, but their revenue per rented hour decreased by 26.2%. This supports that the pricing algorithm optimized for revenue, lowering prices and increasing asset utilization compared to the control group. Providers in T2 experienced slightly lower revenue that providers in T1, but performed better across other metrics. In particular, the revenue per rented hour for providers in T2 decreased by only 12.9%—using their partial pricing control, about 41.8% of the providers in T2 increased the platform-set prices.

To understand the differences in the responses of the providers, we examine the costs incurred by the providers. Providers face various "bring-to-market" (BTM) costs when renting out their assets, which the platform may not be able to observe. These costs can be usefully decomposed into (i) usage-based costs which scale with the rental duration, such as vehicle depreciation, and (ii) transaction-based costs which are incurred per transaction, such as screening the renter, answering questions, and inspecting and cleaning the vehicle after each rental (Filippas et al., 2020). While we cannot observe the BTM costs of the providers directly, we observe proxies for these costs. Our proxy for usage-based BTM costs is the number of miles that providers' vehicles are driven, and our proxy for transaction-based BTM costs is the number of transactions. We find that providers incurred substantially higher BTM costs during the experiment.

Higher costs by themselves do not imply lower utilities because the providers may have been compensated for them by higher revenues. For providers in T1, centralized pricing led to 36.1% lower revenue per mile, and 15% lower revenue per transaction. Providers in T2 saw their BTM cost increase less: their revenue per mile decreased by 23.3%, and their revenue per transaction decreased by 2.5%. The higher costs that treated providers incurred imply that, despite higher revenues, the "true" utility—or profit—from rentals may have decreased. Crucially, only 41.8% of the providers who retained partial control utilized it to increase the platform-set prices by the end of the experiment. This indicates substantial heterogeneity in providers' BTM costs, and that centralized pricing was welfare-increasing, at least for some providers. Regardless, the difference in BTM costs offers a likely explanation for the different responses of the providers in the two treatment groups.

One drawback of the experimental setting is that interference across the experimental units does not allow for direct extrapolation of the experimental estimates of the market effects of centralized pricing to case where prices are set centrally market-wide. To explore this issue, we collected post-experiment data from the platform. Convinced by the findings of the experiment, the platform decided to implement centralized pricing with partial provider control (T2) in the focal market, instead of the originally planned centralized pricing (T1). This market-wide implementation created a quasi-experiment which allows us to examine the market effects of centralized pricing.

We conduct a difference-in-differences analysis using data from the same market during the year before the experiment as the control group. The credibility of our approach depends upon the suitability of the previous year as an appropriate counterfactual; we provide evidence that support the parallel trends assumption, both on the market and the individual level. We obtain estimates that are directionally consistent with the experimental estimates, albeit of smaller magnitude—suggesting that interference plays an important role. The most conservative estimates suggest that centralized pricing with partial provider control increased hours-rented by 24.9% and revenue by 18.7%.

In addition to the economic effects, the market-wide imposition also allows us to examine how providers use the price slider in the long-run. There is substantial heterogeneity in the use of the partial pricing control: while some providers never change the platform-set price, many providers continually experiment, while others set the price slider to its maximum value throughout our data. On average, the increase over the centrally set prices ranged between 3% and 6% over the following one and a half years.

The main contribution of this paper is in offering a detailed account of the challenges of implementing centralized pricing in online marketplaces. We highlight the divergence in platform and provider incentives, by showing that centralized pricing substantially increased revenues, and yet it was with negative reactions, because it did not fully take into account providers' "bring-to-market" costs. Crucially, in the absence of a decentralized price system, market clearing subsequently took place through non-price margins. These non-price adjustments can have pernicious long-term effects for the platform: fewer providers and lower asset availability imply foregone transaction opportunities that might have materialized under provider-set prices, and booking cancelations increase the transaction costs for both sides of the market. We stress that non-price margin adjustments should be taken into account when assessing the welfare effects of market interventions, but that their effects may be hard to quantify both in the short- and in the long-run.

We propose partial control as a way to mitigate the implementation challenges of centralized price-setting. We find that partial control reduced providers' non-price margin adjustments substantially but resulted in similar revenues, and that providers varied greatly in their use of the price slider, indicating substantial preference heterogeneity. Partial control is likely most useful for platforms where providers possess private, heterogeneous, and time-varying information that is economically relevant. In this case, providers can reveal this information without turning to non-price margins, while, crucially, allowing the platform to change the average price level, and increase the intertemporal variation of prices. Importantly, centralized pricing remains nested within the partial control variant; as platforms mature and become better at setting prices, users can choose to relinquish more control to the platform.

We expect our findings to extend to other platforms that are characterized by similar divergence in platform and provider incentives and heterogeneous provider BTM costs. For instance, AirBnB extracts an ad-valorem fee, which incentivizes it to maximize revenue, potentially at the expense of increasing provider (owner) costs. Further, owners incur usage-based BTM costs (e.g., asset wear-and-tear) and transaction-based BTM costs (e.g., finding and dealing with the customer, cleaning the unit, and passing out the keys), which we expect to be heterogeneous across providers. For such a platform, rolling out the partial control feature will allow it to collect private owner BTM cost proxies, so that it can increase owner revenue and asset utilization while mitigating any owner negative reactions.

More broadly, our paper offers partial control as a general product recommendation for online platforms. The partial control feature allows platforms to collect valuable private information from the providers to help improve the platform. For example, Uber has started allowing both drivers and riders to indicate whether they prefer their match to be more or less "talkative"[2], which would allow them to improve upon centralized matching. Similarly, platforms can collect private cost information, which would allow them to implement complex dynamic pricing systems. Traditional dynamic pricing systems need provider costs, which online platforms do not observe. These costs are also difficult to infer because they generally vary significantly across providers and also time. Therefore, a partial control feature is almost necessary to be able to successfully implement such dynamic pricing systems.

Our findings are, of course, subject to some limitations. The most important limitation is that our experiment is conducted on a single platform. As noted above, platforms where suppliers incur similar BTM costs and have access to non-price margins will likely face similar challenges. Future research could verify that providers react similarly to price-setting changes in other marketplaces, as well as changes in the degree of centralization of other types of allocation mechanisms. Another potential limitation is that we examine a single pricing change that decreased average price levels; providers would presumably have responded differently to a "better-designed" pricing system. Regardless, the pricing change increased providers' revenues, and the majority of the providers did not change the centrally set prices; this supports our view that any one price level may be ill-suited for even seemingly similar providers, and that continuously eliciting providers' private information is a key task for centralized allocation mechanisms.

As more of the global economic activity takes place on platforms operating two-sided markets, it becomes increasingly important to recognize and address their unique design challenges. Insofar that these platforms will continue lacking the directive authority that traditional firms typically enjoy over their employees, the users of those platforms will likely

[2]For example, see `https://techcrunch.com/2019/05/14/uber-quiet-ride/`

have access to multiple non-price margins. We document empirically the trade-offs faced by when these platforms make market design interventions, and we propose and evaluate a potential solution.

The rest of the paper is organized as follows. Section 2 surveys related work and introduces the empirical context. Section 3 presents the results of the experiment, and Section 4 the mechanisms behind these results. Section 5 presents quasi-experimental evidence of the market effects of centralized pricing with partial provider control. Section 6 discusses our results, and Section 7 concludes with thoughts on directions for future research.

2 Related work and empirical context

Our study is situated in an online sharing economy platform for short-term vehicle rentals. Sharing economy platforms are online marketplaces that facilitate rentals of durable assets from providers (owners/sellers) to renters (non-owners/buyers). The economic rationale for these markets is that owners of most durable assets use them less than 100% of the time, leaving excess capacity that can be rented out to renters who would like to use the asset (Sundararajan, 2016; Filippas et al., 2020).

Similarly to other online marketplaces, sharing economy platforms extensively use technology to help providers and renters to find, assess, and transact with each other more efficiently than what is possible in most physical markets (Cramer and Krueger, 2016; Einav, Farronato and Levin, 2016; Sundararajan, 2016; Horton, 2017; Filippas, Horton and Golden, 2018; Liu, Brynjolfsson and Dowlatabadi, 2018; Athey and Luca, 2019; Filippas et al., 2020). On the supply side, sharing economy platforms lower the entry costs of smaller providers, affording them substantial flexibility and enabling them to reach buyers more easily (Einav et al., 2016; Chevalier, Chen, Oehlsen and Rossi, 2018). On the demand side, sharing economy platforms expand product variety, with the benefits disproportionately accruing to previous non-owners who gain access to the asset or service (Fraiberger and Sundararajan, 2015; Filippas and Horton, 2018; Filippas et al., 2020).

2.1 Price-setting in online markets

A core market design challenge in online marketplaces is choosing how prices are set. Platforms may choose between several decentralized price-setting mechanisms including posted prices, haggling, auctions, and price recommendations (Einav, Kuchler, Levin and Sundaresan, 2015; Farronato, 2017). However, decentralized price-setting has the potential to lead to inefficient market equilibria, for reasons including buyers' search costs, sellers' price adjustment costs, and sellers' inability to adjust their prices in response to fluctuations in market demand (Diamond, 1971; Li, Moreno and Zhang, 2016; Einav, Farronato, Levin and

Sundaresan, 2018). In light of this, several platforms choose to develop centralized pricing mechanisms—despite the well-known challenges of this undertaking (Hayek, 1945).[3]

Many recent papers examine the implementation of centralized pricing in the context of online marketplaces (Cachon, Daniels and Lobel, 2017; Ma, Fang and Parkes, 2018; Taylor, 2018; Bimpikis, Candogan and Saban, 2019; Gurvich, Lariviere and Moreno, 2019). Yan, Zhu, Korolko and Woodard (2020) provide a comprehensive survey of this literature. The goal of this body of work is to improve the design of centralized pricing; in contrast, we focus on analyzing the effects of centralized pricing on platform users, taking the pricing design as given. A growing number of empirical papers leverage upon variations in the design of centralized pricing. Cohen, Hahn, Hall, Levitt and Metcalfe (2016) use variation in Uber's surge pricing—Uber's response to increases in demand for rides—to estimate the demand elasticity at several points along the demand curve. Building on this methodology, Castillo (2019) estimates a structural model using Uber data, and finds that surge pricing increases welfare, with the benefits disproportionately accruing to the riders. Also in the context of Uber, Hall et al. (2018) find that while changes in fares initially result in higher driver earnings per trip and per hour worked, market adjustments through non-price margins eventually bring driver earnings back to the previous levels.

Our paper contributes to this literature in at least three ways. First, to our knowledge, our paper is the first to examine empirically a market transitioning from decentralized to centralized price-setting—previous work has focused on variations in the design of centralized pricing mechanisms. Second, we examine a sharing economy marketplace where the labor component is less prominent than in ride-hailing marketplaces—but, as we will show, it still plays an important role, affecting providers' costs. Third, we propose and examine empirically a variant of centralized pricing where providers are allowed to retain partial pricing control.

2.2 The focal platform

The setting for our study is a sharing economy marketplace, where providers (owners/sellers) rent out their vehicles to renters (non-owners/buyers). Providers choose when to make their vehicles available, and at what rental rate; renters specify a time period, and perform a map-based online search to book one of the available vehicles. The platform offers typical online marketplace services, including building and maintaining search and reputation systems, curating the matching process, handling payments, and providing insurance and customer support. Crucially, the platform reduces transaction costs by providing proprietary hardware and software for mobile phone-based, keyless unlocking of rented vehicles, as well as parking spots for providers' vehicles.

[3]Even in markets that do not clear on prices, platforms make market design decisions that affect how users "trade" their scarce resources (Filippas and Horton, 2021).

To examine the underlying characteristics of the platform, we begin by presenting some descriptive statistics in Figure 1. On the demand side, most of the rental activity on the platform is short-term—despite the platform allowing rentals lasting up to a month. Panel (a) depicts the histogram of rental durations: about 89 percent of rentals last less than one day, and only 1.6 percent last more than three days. On the supply side, providers vary substantially in how they use the platform. Panel (b) plots the histogram of vehicle availabilities, defined as the percentage of time each vehicle is made available for rentals on the platform each month. Vehicle availabilities vary widely, indicating substantial provider "type" heterogeneity. Regardless, more than 70 percent are "serious" providers, in that they make their vehicles available for rentals at least half of the time—this is consistent with existing empirical evidence suggesting that vehicles are used less than 4 percent of the time.[4]

A distinctive property of the market equilibrium is that a large fraction of the available supply remains unused. Panel (c) depicts the histogram of vehicle utilizations, defined as the percentage of time each vehicle is rented over the percentage of time the same vehicle is made available for rentals. Strikingly, about 65 percent of vehicles have less than 30 percent utilization. Unused capacity may not be a problem in and of itself, or may even be beneficial for the market (Hall, 1983; Castillo, Knoepfle and Weyl, 2017; Hall et al., 2018); that said, the platform employed an ad valorem business model, and believed that increasing asset utilization would boost platform revenues. While the underlying cause for the low utilization rates cannot be known from observational evidence, one potential culprit was the ostensibly suboptimal pricing decisions of the providers. Panel (d) plots the histogram of the average number of price changes made by providers each month. Most providers rarely or never change the rental prices of their vehicles: about 70 percent of the providers changed their price at most once per month. The platform believed that a pricing mechanism change could shift the market to a more efficient equilibrium.

3 Transitioning to centralized pricing

3.1 Experimental design

Motivated by the ostensible link between low utilization and provider pricing, the platform developed a proprietary centralized pricing system. The centralized pricing mechanism attempted to combine data on providers' previous pricing decisions with platform-level supply and demand information to adjust price levels and to increase their intertemporal variation, with the end goal of optimizing revenues. After conducting tests on historical data which yielded positive results, the platform decided to evaluate the new pricing system within the marketplace. Toward that end, the platform introduced two versions of centralized pricing

[4] See `https://www.racfoundation.org/research/mobility/spaced-out-perspectives-on-parking`

Figure 1: Descriptive statistics of the rental activity in the focal platform

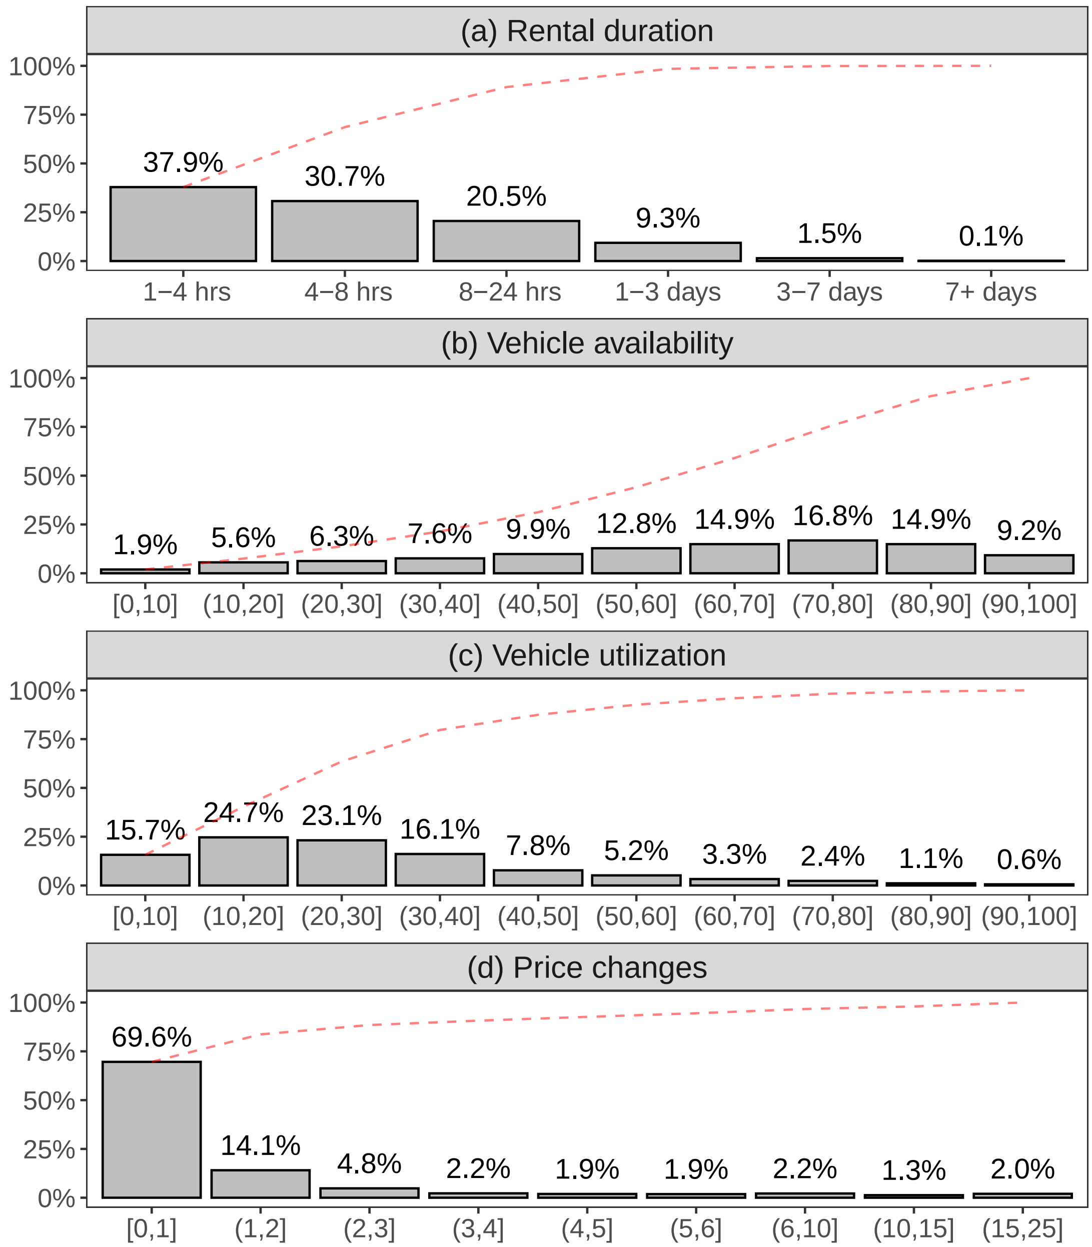

Notes: This figure reports descriptive statistics for rentals in the focal platform. Panel (a) plots the distribution of rental durations. Panel (b) plots the distribution of monthly vehicle availability, defined as the percentage of time that a vehicle was made available for rentals. Panel (c) plots the distribution of monthly vehicle utilization, defined as the percentage of time that a vehicle was rented out over the percentage of time that it was made available. Panel (d) plots the distribution of provider price changes per month. For each panel, the values are discretized into categories, the value of each bin is shown above it, and the red dashed line depicts the corresponding cumulative distribution function.

experimentally in one if its largest and most mature markets.

The experiment aimed to evaluate two centralized pricing variants. In the first version of the treatment (T1), the platform assumed complete pricing control. In the second version of the treatment (T2), the platform assumed partial control: whereas prices and intertemporal price variation were determined and implemented by the platform's centralized pricing mechanism, providers were able to raise or lower the centrally set prices by up to 30 percent. The rationale behind this "hybrid" approach was that providers could use the partial price-setting control to reveal their idiosyncratic preferences. A control group (T0) remained at the status-quo pricing system, retaining full price-setting control.

Eligible providers were those who had made their vehicles available for at least 24 hours during the month prior to start of the experiment. Amongst those providers, the platform selected a random subset to be included in the experiment, with the final sample consisting of 1,218 subjects. Each subject was randomly assigned to one of the two treatment groups with probability 13.5% each, or to the control group with probability 73%. The experiment lasted for about two months; the size and length of the experiment was determined by an ex ante power calculation conducted by the platform.[5]

In Appendix A, we report statistical tests for pairwise mean comparisons across several pre-randomization attributes, and find that the randomization was successful. We also report the pricing interfaces, along with all information that subjects had access to. Furthermore, we found no evidence that subjects communicated with each other during the experiment.[6]

3.2 Provider responses

We begin by examining the responses of treated providers to the pricing change. Throughout the rest of this section, our estimation strategy is to regress each outcome of interest on treatment indicators, that is,

$$y_j = \beta_0 + \beta_1 \mathrm{T}_{1j} + \beta_2 \mathrm{T}_{2j} + \epsilon, \tag{1}$$

where y_j is the outcome of interest, T_{ij} is an indicator variable for whether subject j was assigned to treatment i, and ϵ is an error term. We report the estimated average treatment effect in Figure 2i by plotting the least squares estimates $\hat{\beta}_i$ for each of the two active treatment groups, along with a 95% confidence interval around each point estimate. All regression

[5]The platform intended to conduct an experiment with sufficient power to detect a 5 percentage point change in the probability of a user dropout, at 90% power. As we show in what follows, the "realized" power for the main experimental outcome was close to 100%. We do not report the fraction of users that were included in the experiment for confidentiality purposes.

[6]To be sure, it is not possible rule out completely the possibility that subjects communicated with each other. As an indirect test, we show in Appendix A that control group providers did not change the prices of their vehicles more frequently during the experimental period—we would expect untreated providers to have done so upon finding out that other providers have become more price-competitive.

tables are reported in Appendix B.

The imposition of the centralized pricing mechanism had substantial effect on the exit decisions of the treated providers—which was also the primary outcome of interest for the platform. Panel (a) reports estimates from regressions of provider exit on the treatment indicators. The exit rate for the treatment group T1 increased by 30.5 percentage points from a baseline exit rate of 8.9 percentage points for T0. Allowing providers to retain some pricing control substantially reduced this effect: the exit rate in the T2 group increased by about 15 percentage points, which was 50.7% lower than the T1 group.[7]

Providers also reduced the availability of their vehicles in response to the new pricing mechanism. Panel (b) reports estimates from regressions of vehicle availability on the treatment indicators; availability is defined as the percentage of time each provider makes their available on the platform, ranging from 0 for never available, to 1 for always available. The reported estimate of the average treatment effect includes all treated providers.[8] Providers in the T1 and T2 groups reduced their vehicle's availability by about 18.9 and 12.2 percentage points respectively. Allowing providers to retain some pricing control substantially ameliorated availability reductions: the availability reduction for providers in T1 was 54.7% higher than that of providers in T2.

Providers also used transaction cancelations as an additional lever of control. Panel (c) reports estimates from regressions of provider cancelation rates on the treatment group indicators; cancelation rate is defined as the percentage of transactions canceled by each provider, ranging from 0 for providers who did not cancel any transaction, to 1 for providers who canceled every transaction. The sample includes providers who exited during the experimental period, setting their cancelation rate to one for all periods after their time of exit. Providers in treatment group T1 increased their cancelation rate by about 19.8 percentage points, and providers in treatment group T2 increased their cancelation rate by about 12.6 percentage points. Again, the difference between the cancelation responses between the two treatment groups is substantial: the cancelation in T2 was 36.4% lower than in T1.

3.3 Economic outcomes

We next examine the experimental effects on providers' core economic outcomes. We use the estimation strategy described in Equation (1). In Figure 2ii, we report the estimated effects as percentage changes over the control group outcomes, by plotting the least squares

[7]Note that here—and throughout the paper—for differences in levels where the outcome is naturally discussed as a fraction, we label level differences as "percentage points." For percentage changes with respect to the outcome of another experimental group, we use the "%" symbol.

[8]We find a similar patterns of results if we restrict our sample to providers who did not exit during the experiment; we report these analyses in Appendix B. In our context, this restriction is likely to yield conservative estimates of the average treatment effect on provider availability and cancelation rates: providers who exited would have presumably exhibited the most pronounced responses (Gerber and Green, 2012).

Figure 2: Experimental estimates of the treatment effects on outcome variables

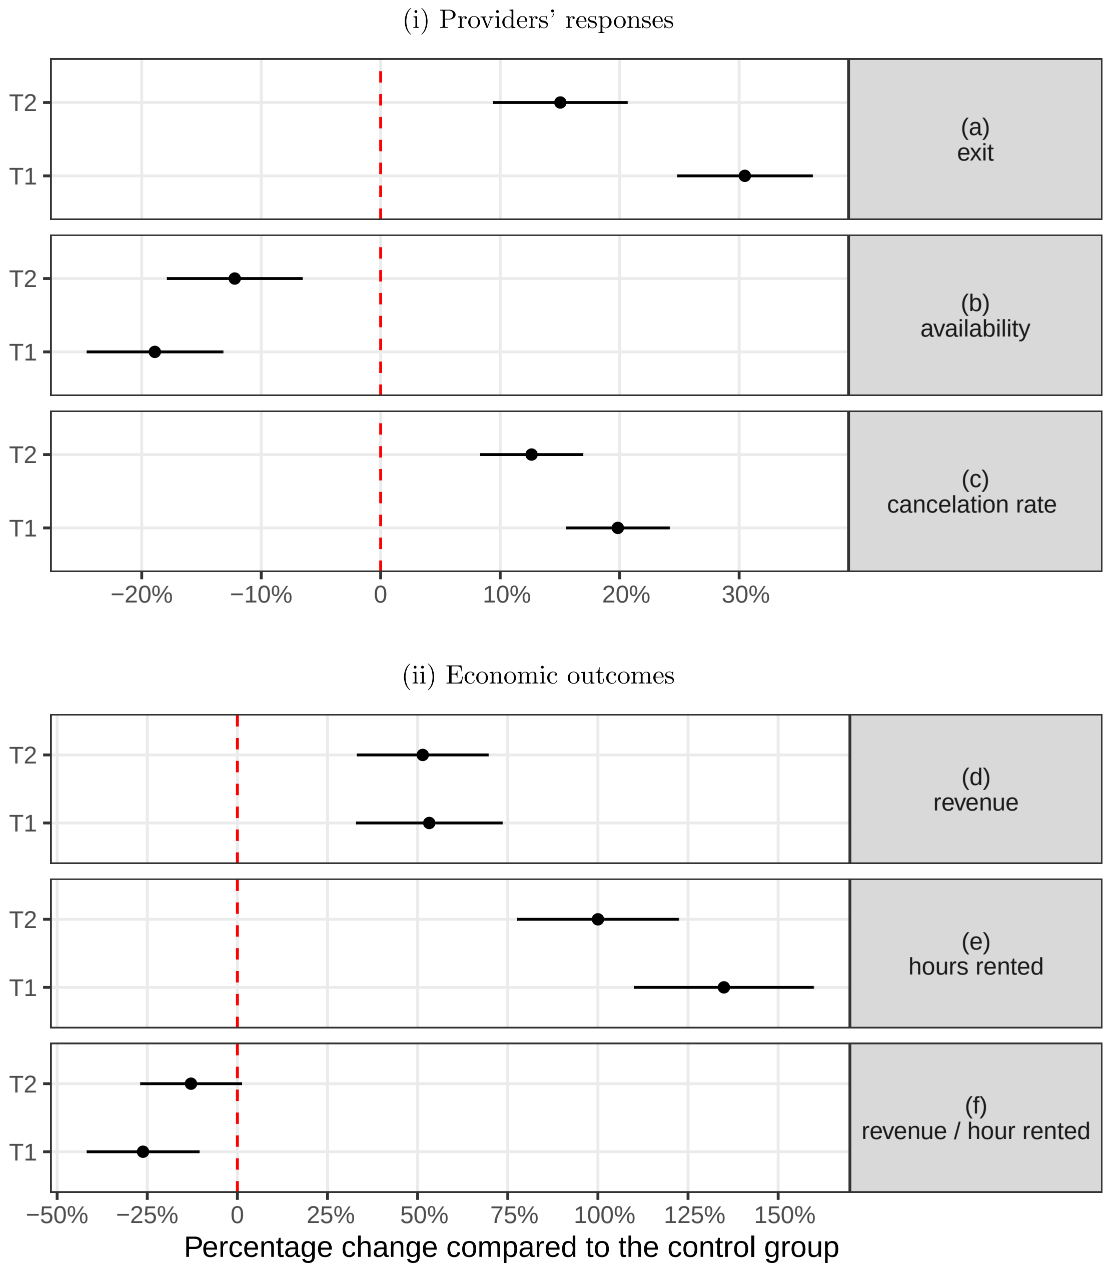

Notes: This figure reports estimates of treatment effects on providers' responses and economic outcomes. Each panel plots the percentage change in the dependent variable versus the control group for the treatment groups. The dependent variables are calculated for the experimental period, and are (a) provider exit from the platform, (b) vehicle availability, (c) transaction cancelation rate, (d) provider revenue, (e) vehicle hours-rented, and (f) revenue per hour rented. A 95% confidence interval is plotted around each estimate. All regression tables, estimates employing alternative samples, and alternative representations of the results can be found in Appendix A and Appendix B.

estimates $\hat{\beta}_i/\hat{\beta}_0$ for each of the two active treatment groups, along with a 95% confidence interval around each point estimate. Our sample comprises providers who did not exit the platform throughout the experiment.[9] All regression tables are reported in Appendix B.

The centralized pricing mechanism increased revenue for treated providers. Panel (d) reports estimates from regressions of provider revenue on the treatment indicators. Providers in T1 saw their revenue increase by 53.2%, and providers in T2 saw their revenue increase by 51.4%, compared to the average revenue of providers in the control group. This revenue increase was a consequence of an increase in the intensive margin of rentals. Panel (e) reports estimates from regressions of the number of hours vehicles were rented on the treatment indicators. Compared to the control group, vehicles of providers in T1 were rented out 135% more, and vehicles of providers in T2 were rented out 100% more. Together, these results indicate that the pricing algorithm achieved its objective of maximizing provider revenues.

The vehicles of treated providers were priced lower by the centralized pricing algorithm. Panel (f) reports estimates from regressions of the revenue per rented hour on the treatment indicators. For every hour their vehicles were rented out, providers in treatment group T1 made 26.2% less revenue than control group providers. The same decrease was only 12.9% for providers in treatment group T2; about half of these providers utilized their partial pricing control to increase prices for their vehicles (see also the discussion in Section 6).

Taken together, the experimental findings presented in this section provide evidence that allowing providers to retain partial control ameliorated their negative non-price responses substantially, at almost no cost—providers in the two treatment groups experienced similar revenue increases. It is worth noting that interference across subjects does not allow us to extrapolate the experimental estimates for revenue and hours-rented to the case where prices are set centrally for every provider. Section 5 provides quasi-experimental estimates of the market-wide effects of centralized pricing.

4 Reasons for the providers' responses

Centralized pricing increased the revenue of treated providers by increasing the utilization of their assets. However, higher provider revenues need not necessarily imply higher provider utilities: the reason is that providers incur "bring-to-market" (BTM) costs when renting out their assets (Filippas et al., 2020).

BTM costs can be decomposed into (i) *usage-based* BTM costs, which are analogous to duration of the rental, and include labor costs, asset depreciation, and complementary consumables, and (ii) *transaction-based* BTM costs, which providers incur on a per-rental

[9] We construct the sample in that manner in order to not bias the estimates of the average treatment effect in favor of the treatment groups that experienced lower exit rates. We find similar patterns of results using alternative samples for each outcome variable. We report these analyses in Appendix A.

basis, and include the cost of finding a trading partner, coming to terms, executing payments, and handing off and "resetting" the asset.[10] In our empirical context, the main component of usage-based BTM costs is vehicle depreciation through mileage increases, and transaction-based costs include screening the renter, answering questions, and inspecting and cleaning the vehicle after each rental. We cannot hope to account for all potential sources of BTM costs, as these costs are not directly observable—this is precisely one of the challenges in implementing centralized pricing. Instead, we sidestep this issue by examining how the pricing mechanism change affected proxies for these costs.

Our proxy for usage-based BTM costs is the number of miles driven by each vehicle and our proxy for transaction-based BTM costs is the number of transactions made on the platform; more miles driven and more transactions indicate higher usage- and transaction-based BTM costs, respectively. In Figure 3, we report the estimated effects as percentage changes, by plotting the least squares estimates $\hat{\beta}_i/\hat{\beta}_0$ for each of the two active treatment groups, along with a 95% confidence interval around each point estimate. We use the estimation strategy described in Equation (1), with miles driven and transactions as the outcome variables, and our sample comprising providers who did not exit the platform throughout the experiment.[11] All regression tables are reported in Appendix B. Centralized pricing increased BTM costs for both treatment groups, with the change being more pronounced for T1.

Higher costs by themselves do not imply lower utilities because the providers may have been compensated for them by higher revenues. While we cannot ascertain this for sure, we can compute proxies to assess the impact of centralization on provider margins. To isolate the effects of the two types of costs, we compute proxies for the usage- and transaction-based margins. Our proxy for the usage-based margin is the revenue per mile each vehicle was driven, and our proxy for the transaction-based margin is the revenue per transaction made on the platform. We follow the same estimation strategy as above, and we report the estimated effects as percentage changes in Figure 3.

The centralized pricing mechanism decreased substantially providers' revenue per mile. Compared to the control group, revenue per mile decreased by 36.1% for providers in T1. These effects were less pronounced for providers in T2, who saw their revenue per mile decrease by 23.3%. It is worth noting that revenue per mile did not fall below $0.30/mile for any transaction in our data—this is the median cost-per-mile estimate for vehicle operating expenses obtained by Zoepf, Chen, Adu and Pozo (2018). Treated providers also experienced a decrease in their revenue per transaction. More specifically, providers in T1 experienced

[10]For example, the usage-based BTM costs of driving with Uber include the driver's labor, the increases in the car's mileage, and the gas consumed while driving. The transaction-based BTM costs include the costs of finding the passenger, and verifying her identity. Online platforms typically reduce these costs by taking advantage of technological advances, including internet- and GPS-enabled smartphones equipped with digital cameras (Varian, 2010; Filippas et al., 2020).

[11]Alternative estimation strategies yield similar patterns of results; see Appendix A for more details.

Figure 3: Experiment estimates of the treatment effects on BTM costs and margins.

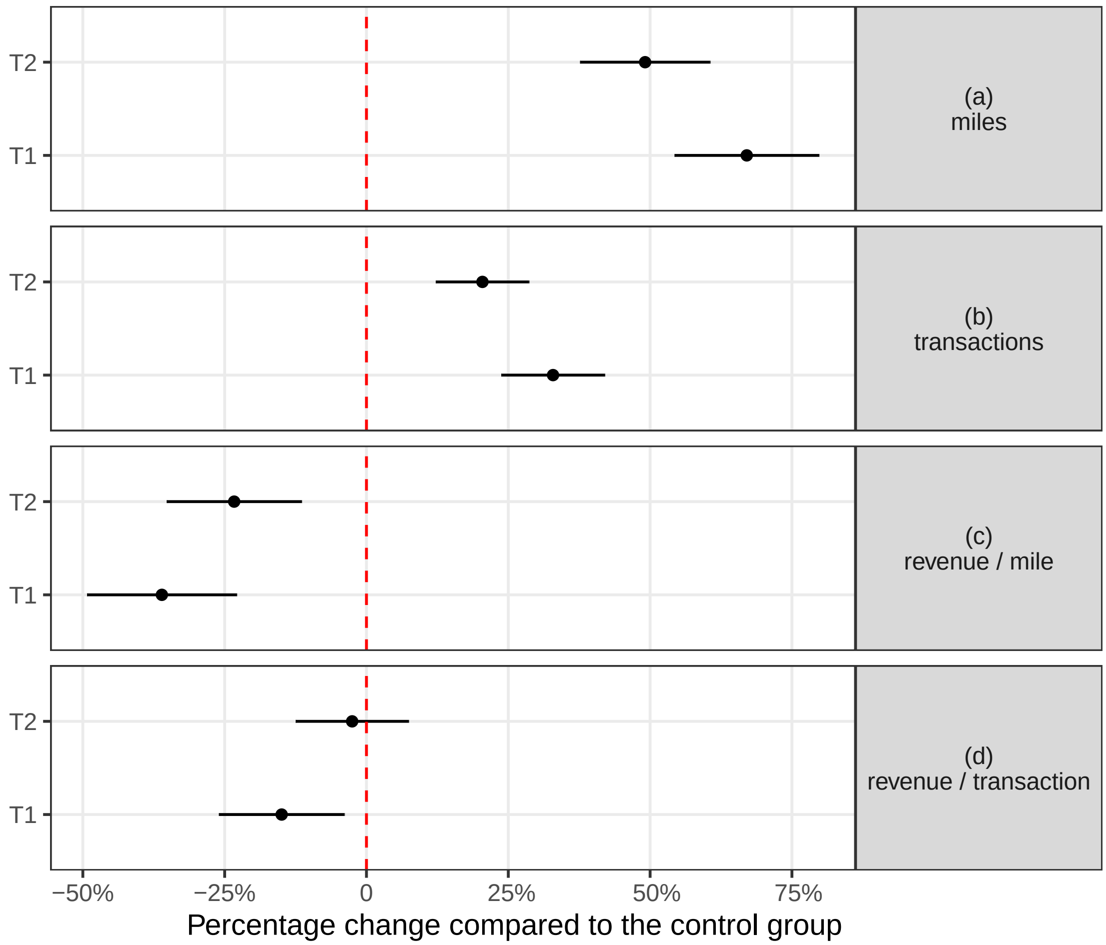

Notes: This figure reports estimates of treatment effects on proxies for providers' BTM costs and margins. Each panel plots the percentage change in the dependent variable versus the control group for the treatment groups. The dependent variables are calculated for the experimental period, and are (a) miles driven, (b) completed transactions, (c) revenue per mile driven, and (d) revenue per completed transaction. A 95% confidence interval is plotted around each estimate. All regression tables, estimates employing alternative samples, and alternative representations of the results can be found in Appendix A and Appendix B.

a 15% decrease compared to control group providers, whereas providers in T2 experienced only a 2.5% decrease—the latter estimate is statistically indistinguishable from zero.

The experimental results imply that revenues are not increasing in proportion to the costs, putting pressure on provider margins. This effect is particularly pronounced for the providers in T1, offering an explanation for their stronger reactions to the centralization of price setting. This finding also highlights the challenges with centralized pricing that are unique to market-based businesses but absent from traditional revenue management systems: when platform users have heterogeneous preferences, partial control may be preferable to complete centralization.

An alternative hypothesis that could explain our findings posits that providers value having control over their assets for non-economic reasons. For example, if providers are averse to changes in what was "implicitly agreed" upon joining the platform, they may exhibit more pronounced negative responses when they lose "more" control. Although we cannot distinguish between this explanation and the BTM-cost-based explanation directly, we can provide an indirect test by examining the responses in provider subpopulations with distinctly different platform tenures. We find that providers with different tenures exhibit similar responses to the pricing change, suggesting that familiarity or inertia has little impact on provider behavior—more details are provided in Appendix A. This supports our view that providers in online markets care about their relative costs and benefits, and hence non-economic considerations are likely second-order.

5 Effects of market-wide centralization of price-setting

The pricing experiment allows us to examine how and why providers react to centralized price-setting, as well as whether partial provider control can change how providers respond. At the same time, the interference between the experimental units limits our ability to extrapolate the experimental estimates of the economic effects of centralized pricing to the case where prices are set centrally for every provider in the market.

After the experiment concluded, the platform implemented centralized pricing with user control (T2) in the focal market of our study. This market-wide implementation resulted in a quasi-experiment that allows us to obtain estimates of the market-wide effects of centralized pricing by comparing market outcomes before and after the pricing change. To control for seasonal differences, we adopt a difference-in-differences estimation strategy using data for the same market from one calendar year prior to the platform intervention.

The credibility of our difference-in-differences approach depends on the suitability of the prior year as an appropriate counterfactual. To assess this, Figure 4i juxtaposes the weekly platform revenue and hours-rented time series for the actual and the prior years. The outcomes for the actual year are depicted using a solid line, and the outcomes for the prior year are depicted using a dashed line. Week "0" is the week when the experiment commenced, and week "8" is the week when the experiment concluded and centralized pricing was imposed market-wide during the actual year. The two time series trends track closely in the pre-period, but then diverge in the post-period. The post-period difference seems to be caused mainly by a decrease in the prior year that is not matched in the actual year. For robustness, we include additional analyses of the parallel trends assumption in Appendix A.5.

To be sure, we cannot directly rule out the hypothesis that the observed divergence may have been caused by factors unrelated to the centralization of price-setting, in either or even

Figure 4: Market effects of centralized pricing with user control

(i) Longitudinal representation

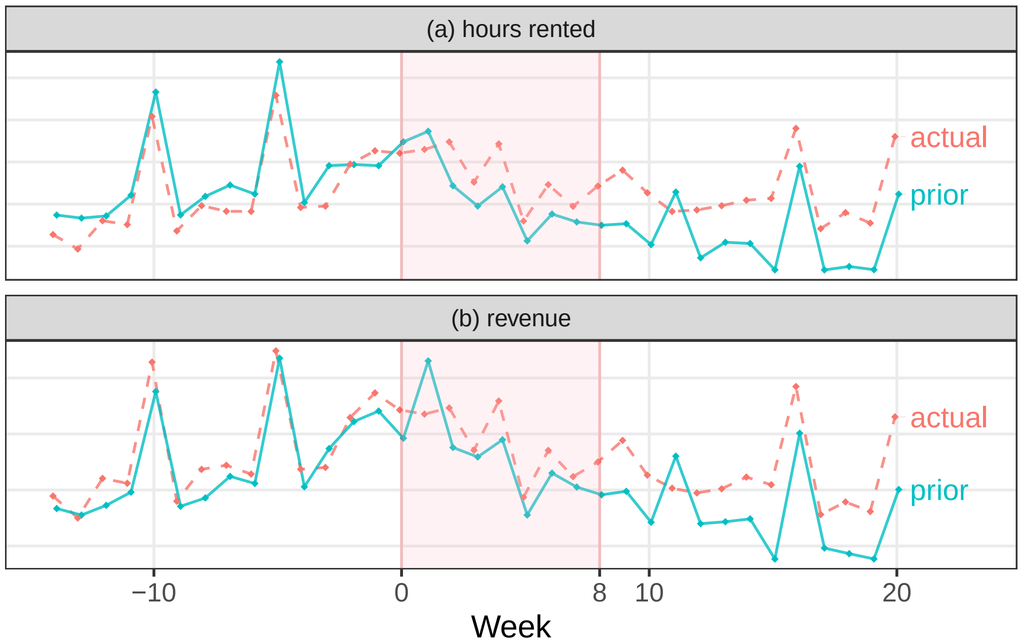

(ii) Difference-in-difference estimates

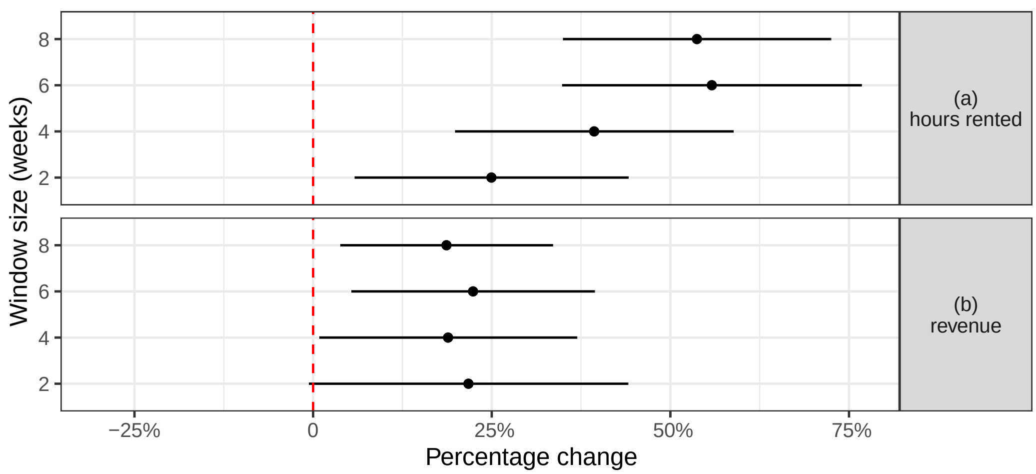

Notes: This figure reports the effects of centralized pricing with partial user control. The top panel plots the platform hours-rented and revenue for each week in our data. The pink-shaded area indicates the experimental period; on its left, prices were set by the providers (T0), and on its right, prices were set centrally with partial provider control (T2). The bottom panel reports difference-in-differences estimates, calculated as the actual year estimate over the estimate from the prior year. The dependent variables are calculated for different pre- and post-period windows around the event, indicated on the y-axis, effects are plotted as the percentage change in the dependent variable, and a 95% confidence interval is plotted around each estimate. All regression tables are reported in Appendix B.

both years. To our knowledge, no such change took place in the marketplace, and we were unable to find evidence of any such change in our data. Furthermore, we observe that the divergence between the two time series is (i) smaller between week "0" and week "8," when price-setting was centralized only for providers participating in the experiment during the actual year, (ii) larger after week "8" when centralized pricing was imposed market-wide, and (iii) otherwise exhibits a functionally similar trend across both years. It is worth noting that the two time series are similar in magnitude in the pre-period, indicating that a similar "amount" of economic activity has been taking place in the focal market over time. This is not due to chance: the focal market is one of the most mature markets of the platform. Together, these observations provide strong evidence in support of the suitability of the prior year as an appropriate counterfactual.

Our empirical strategy is to estimate a regression of the form

$$y_j = \beta_0 + \beta_1 \textsc{Post}_j + \beta_2 \textsc{Trt}_j + \beta_3 \textsc{Post}_j \times \textsc{Trt}_j + \epsilon, \tag{2}$$

where $\textsc{Post}_j$ is a variable indicating whether the outcome is measured after the time of the market-wide imposition in either year, $\textsc{Trt}_j$ is a variable indicating whether the outcome is measured during the "actual" year, and $\textsc{Post}_j \times \textsc{Trt}_j$ is the interaction of the two indicators. The implied difference-in-differences treatment effect estimate is $\hat{\beta}_3$.

We provide estimates using revenue and hours-rented aggregates during the pre- and post-intervention windows, discarding observations during the experimental period for both years. Essentially, we "collapse" the time series into a pre- and post-period (Bertrand, Duflo and Mullainathan, 2004). As the choice of pre- and post-period window size may affect our estimates, we simply derive estimates using various window size values. We report the difference-in-differences estimates as percentage changes in Figure 4ii, and we report all regression tables in Appendix B.

Starting with revenue, all estimates are positive, and range from 18.6% to 22.3%. The estimate using the shortest time window is not conventionally statistically significant, but with larger post-period windows the estimates become "more" statistically significant. This pattern is explained by the fact that some transactions taking place in the post-period were booked before price-setting was centralized; as we increase the window size, the relative effect of these transactions on the estimates diminishes.

With regards to hours-rented, all estimates are positive and highly statistically significant, and range from 24.9% to 55.7%. The estimate with the smallest magnitude is obtained when we use the shortest post-period window. It is worth noting that, taken together, the estimates imply that the revenue per hour rented on the platform decreased following the centralization of price-setting.

The difference-in-differences estimates are directionally consistent with the experimental estimates, and indicate that centralized pricing has substantial market effects. However, the magnitude of the difference-in-differences estimates is smaller. Two broad—and not mutually exclusive—sets of reasons may have led to this divergence: (i) that allowing *all* providers to retain partial pricing control resulted in lower revenue and lower utilization because providers increased the centrally set price more than they did during the experiment, and (ii) that the market effects of centralized pricing are smaller than the estimates obtained through the experiment because of interference. While we cannot distinguish between these two sets of reasons directly, we provide evidence in Section 6 that providers increased the centrally set prices, albeit not substantially more than during the experiment. This suggests that interference likely explains much of the observed difference between the experimental and the quasi-experimental estimates.

6 Discussion

6.1 Partial control as a market design feature

Providers reacted to the centralization of price-setting by exiting the platform, reducing the availability of their assets, and canceling booked transactions. Market clearing through these non-price margins can be particularly detrimental to the platform. For example, fewer providers and lower asset availability result in foregone transaction opportunities that might have otherwise materialized at the provider-set prices. Furthermore, transaction cancelations increase the transaction costs for both sides of the market, and may erode the users' trust in the reliability of the platform, thereby causing further exit. Partial provider control ameliorated this problem considerably: as we showed, providers who retained partial price-setting control reacted substantially less through non-price margins.

Provider control also creates a channel through which the platform can obtain economically relevant information. Although modern platforms rely overwhelmingly on technological solutions to obtain granular market-wide information, idiosyncratic and time-varying private information that providers possess may remain hard-to-access.[12] The hardness of eliciting private information is a recurrent criticism of centralized allocation mechanisms that dates at least back to Hayek (1945). Provider control functions as a channel through which providers reveal their private information as it changes, and can help the platform to circumvent the information elicitation challenge.

Crucially, the "fully" centralized version of the pricing mechanism remains nested within the partial control version—should a provider choose to accept the centrally set price as

[12] For example, a provider whose parents are visiting for the weekend may experience a positive shock in her reservation price, as her utility from using her vehicle increases. On a sunny day, some providers may have higher reservation prices, whereas providers who are particularly prone to sunburns and photic sneezes may have lower reservation prices.

is. As platforms mature, collect data, conduct experiments, and increase their technological sophistication, presumably they become better at setting prices. Subsequently, providers may choose to relinquish more control to the platform, without the need for a market intervention. Conversely, degradations in the quality of the centralized pricing system, e.g., due to an implementation error, can be "signaled" through the partial control channel rather than through much costlier non-price adjustments.

We laid out several reasons why imbuing centralized market mechanisms with some degree of provider control can be a valuable tool in the hands of market designers. It is also worth examining how providers utilized the partial price-setting control feature in our context. Figure 5i plots the distribution of price slider levels for providers in the experimental group T2 at the end of the experimental period—recall that these providers could adjust the centrally set price by up to 30 percent. About 54.6% of the providers did not utilize the slider, while most of the other providers increased the centrally set prices. This variation in the use of the price slider provides some evidence of substantial preference heterogeneity.

Figure 5ii plots the providers' average price slider level over time after the market-wide price centralization. The average slider level remains close to 104% initially, and increases by about 2 percentage points during month "9." Of course, we cannot know whether this jump is due to a shift in the providers' BTM costs, a change in the market demand for rentals, a new marketplace feature, or even due to chance. Regardless, partial control allows providers to respond to this ostensible change, even if the platform does not.

6.2 Welfare effects of centralized price-setting

The market equilibrium following the centralization of price-setting was characterized by lower prices, higher utilization, and higher platform revenue. Lower prices likely increased renter surplus, at least insofar as they outweighed the impact of the providers' non-price responses. This finding is in congruence with previous results indicating that centralized pricing increases consumer-side surplus (Castillo et al., 2017; Dinerstein, Einav, Levin and Sundaresan, 2018; Castillo, 2019).

On the provider side, the effects of centralized pricing are less clear. While the revenues of the providers increased, so did the costs of renting out assets on the platform. As such, the overall effect likely varies at the individual provider level; the differences in how providers reacted to the pricing change, as well as in how they utilized the price slider provides some evidence in support of this claim. Nevertheless, the non-price responses we documented in Section 4 indicate that the effect was negative, at least for some providers.

There are likely different welfare-maximizing price-setting mechanisms for differently organized markets, depending on the particular problems that each market has to address (Castillo et al., 2017). In our setting, we showed that centralized pricing increases platform revenue,

Figure 5: Short- and long-run use of the price slider.

(i) Distribution of price slider levels at the end of the experimental period

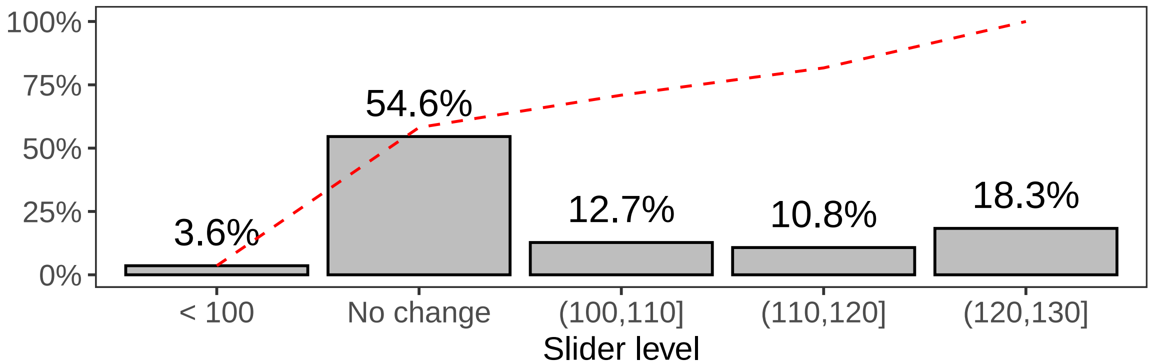

(ii) Average price slider level over time after the experiment

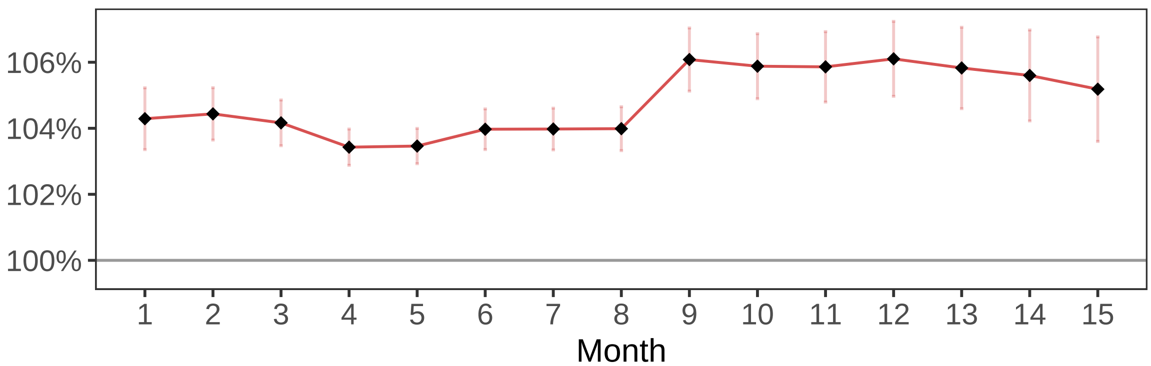

Notes: This figure shows how providers use the slider in the short- and in the long-run. The top panel plots the distribution of price slider choices at the end of the experimental period for providers in the experimental group T2, with slider levels discretized into 5 categories. The value of each bin is shown above it, and the red line depicts the cumulative distribution function. The bottom panel plots the mean slider level over time after the end of the experimental period. The average slider level is computed for every month, and a 95% confidence interval is shown for each mean.

but also results in providers reacting through non-price margins. As we reasoned, an important shortcoming of non-price responses is that they can affect user retainment in the long-run, both on the provider and on the renter side. One solution to this problem is allowing providers to retain partial pricing control. The optimal degree of provider control likely depends on the relative costs and benefits of price and non-price responses, as well as on other characteristics of each market (Farronato, 2017; Hall et al., 2018).

6.3 External validity

Our analysis focuses on a single sharing economy platform. It is worth thinking through how our results may generalize to other markets. The structure of the costs and benefits providers incur in other marketplaces seems to be similar to the focal market. For example, ride-hailing platform Uber extracts an ad-valorem fee, and providers (drivers) incur usage-based BTM costs including labor costs, car mileage increases, and gas costs, as well as transaction-based BTM costs including finding and picking up passengers, and verifying their identity. Similarly, home-sharing platform Airbnb also extracts an ad-valorem fee, and providers (owners) incur usage-based BTM costs including asset wear-and-tear, as well as transaction-based BTM costs including finding and dealing with the customer, cleaning the unit, and passing out the keys. Mirroring our findings, Hall et al. (2018) show that market-clearing can take place through non-price margins on Uber—which also employs centralized pricing—and casual empiricism further supports this claim.[13] As such, providers' decisions in other marketplaces are likely characterized by similar economics, and hence we expect the findings of our study to generalize reasonably well beyond the marketplace we studied.

We found that partial provider control can be a beneficial addition to centralized pricing. The usefulness of embedding provider input into centralized allocation mechanisms may extend well beyond the case of price-setting. For example, Uber takes on the problem of matching drivers with passengers centrally, presumably because both sides are characterized by little preference heterogeneity. Recently, Uber has started allowing both drivers and riders to indicate whether they prefer their match to be more or less "talkative."[14] This match-relevant "type" information is private, presumably time-varying, and can be utilized to improve upon the centralized matching. Despite its promising outlook, whether market designers can usefully embed partial provider control in non-price centralized allocation mechanisms is a question that ultimately needs to be answered empirically.

7 Conclusion

Summary

This paper provides experimental and quasi-experimental evidence of the effects of centralized price-setting in an online marketplace. The key finding is that centralized pricing resulted in an equilibrium characterized by lower prices and higher utilization that benefited renters and the platform, at least in the short-run. However, centralized pricing had ambiguous effects for the providers because it increased their transaction costs, and providers responded

[13] Uber drivers sometimes cancel trips when they are not satisfied by the platform-set price and match. For some anecdotal accounts, see `https://qz.com/1387942/uber-drivers-are-forcing-riders-to-cancel-trips-when-fares-are-too-cheap`.

[14] For example, see `https://techcrunch.com/2019/05/14/uber-quiet-ride/`

through non-price adjustments with potentially deleterious consequences. Allowing providers to retain partial pricing control resulted in substantially less negative responses but similar efficiency gains.

Insights for platform designers

Our work provides several lessons for platform designers. First, it highlights the key tradeoffs that a platform faces when centralizing important provider decisions. On the one hand, centralization gives more control to the platform so it can offer a more consistent customer experience, and compete better with traditional firms offering similar services. On the other hand, centralization can result in unintended and potentially unanticipated negative reactions from the providers. A platform designer must account for such reactions carefully in making centralization decisions.

Second, our work sheds light on two root causes for provider reactions: (a) the existence of provider private information, such as information about costs, which is unknown to the platform but relevant for making centralization decisions, and (b) a misalignment of incentives between the platform and the provider, even with an ad-valorem fee structure. The lack of private provider information limits the platform's ability to make effective decisions, particularly when provider preferences are sufficiently heterogeneous and time-varying. In such cases, it is essential for the platform to devise mechanisms to elicit private information, the absence of which can also lead to mistrust among the providers. Even with access to such private information, the platform faces an incentive misalignment issue. A platform's desire to maximize its revenues and ensure a consistent customer experience often conflicts with the need to account properly for provider-side costs and preference heterogeneity. While such misalignments cannot be fully resolved, eliciting private provider information might allow the platform designer to strike a better balance between these conflicting objectives.

Third, our work shows that an in-between solution of allowing the providers retain some control can result in significantly better provider reactions. This is particularly true when the platform starts with a decentralized design and then decides to transition to a centralized design, which is quite common in practice. Such improvement in provider reactions can even come with little to no impact on the overall performance—we find this to be the case in our study, where several providers did not exert the control provided to them by the platform. Platform designers should explore partial control designs before moving to complete centralization.

Future work

A natural direction for future work would be to replicate the results of this study, particularly in other online marketplaces, using recently developed experimental design techniques (Johari, Li, Liskovich and Weintraub, 2020), and utilizing larger samples. Finding ways to embed partial control in market mechanisms beyond price-setting is another interesting next step. Future research could also examine how providers who retain partial control use it to respond to shifts in their utility functions, whether they relinquish more control to the platform over time, as well as their rate of learning. A question left unanswered is determining the optimal "amount" of control that the platform should relinquish to its users.

References

Angrist, Joshua D and Alan B Krueger, "Empirical strategies in labor economics," in "Handbook of labor economics," Vol. 3, Elsevier, 1999, pp. 1277–1366.

Athey, Susan and Michael Luca, "Economists (and economics) in tech companies," *Journal of Economic Perspectives*, 2019, *33* (1), 209–30.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How much should we trust differences-in-differences estimates?," *The Quarterly journal of economics*, 2004, *119* (1), 249–275.

Bimpikis, Kostas, Ozan Candogan, and Daniela Saban, "Spatial pricing in ride-sharing networks," *Operations Research*, 2019.

Cachon, Gerard P, Kaitlin M Daniels, and Ruben Lobel, "The role of surge pricing on a service platform with self-scheduling capacity," *Manufacturing & Service Operations Management*, 2017, *19* (3), 368–384.

Castillo, Juan Camilo, "Who Benefits from Surge Pricing?," *Available at SSRN 3245533*, 2019.

_ , Dan Knoepfle, and Glen Weyl, "Surge pricing solves the wild goose chase," in "Proceedings of the 2017 ACM Conference on Economics and Computation" ACM 2017, pp. 241–242.

Chevalier, Judith A, M Keith Chen, Emily Oehlsen, and Peter E Rossi, "The value of flexible work: Evidence from Uber drivers," *Journal of Political Economy*, 2018.

Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, "Using big data to estimate consumer surplus: The case of uber," Technical Report, National Bureau of Economic Research 2016.

Cramer, Judd and Alan B Krueger, "Disruptive change in the taxi business: The case of Uber," *The American Economic Review*, 2016, *106* (5), 177–182.

Diamond, Peter A, "A model of price adjustment," *Journal of economic theory*, 1971, *3* (2), 156–168.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan, "Consumer price search and platform design in internet commerce," *American Economic Review*, 2018, *108* (7), 1820–59.

Einav, Liran, Chiara Farronato, and Jonathan Levin, "Peer-to-peer markets," *Annual Review of Economics*, 2016, *8*, 615–635.

_ , _ , _ , and Neel Sundaresan, "Auctions versus posted prices in online markets," *Journal of Political Economy*, 2018, *126* (1), 178–215.

_ , Theresa Kuchler, Jonathan Levin, and Neel Sundaresan, "Assessing sale strategies in online markets using matched listings," *American Economic Journal: Microeconomics*, 2015, *7* (2), 215–47.

Farronato, Chiara, "Pricing mechanisms in online markets," *Working Paper*, 2017.

Filippas, Apostolos and John J Horton, "The tragedy of your upstairs neighbors: When is the home-Sharing externality internalized?," *Working paper*, 2018.

_ and _ , "The Production and Consumption of Social Media," Technical Report, National Bureau of Economic Research 2021.

_ , _ , and Joseph Golden, "Reputation inflation," *Working Paper*, 2018.

_ , _ , and Richard J Zeckhauser, "Owning, using and renting: Some simple economics of the "sharing economy"," *Management Science*, 2020.

Fraiberger, Samuel P and Arun Sundararajan, "Peer-to-peer rental markets in the sharing economy," *Working Paper*, 2015.

Gerber, Alan S and Donald P Green, *Field experiments: Design, analysis, and interpretation*, WW Norton, 2012.

Gurvich, Itai, Martin Lariviere, and Antonio Moreno, "Operations in the on-demand economy: Staffing services with self-scheduling capacity," in "Sharing Economy," Springer, 2019, pp. 249–278.

Hall, Jonathan V, John J Horton, and Daniel T Knoepfle, "Pricing efficiently in designed markets: Evidence from ride-sharing," 2018.

Hall, Robert E, "Is Unemployment a Macroeconomic Problem?," *The American Economic Review*, 1983, *73* (2), 219–222.

Hayek, Friedrich August, "The use of knowledge in society," *The American economic review*, 1945, *35* (4), 519–530.

Hirschman, Albert O., *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*, Vol. 25, Harvard university press, 1970.

Horton, John J, "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

Johari, Ramesh, Hannah Li, Inessa Liskovich, and Gabriel Weintraub, "Experimental Design in Two-Sided Platforms: An Analysis of Bias," *arXiv preprint arXiv:2002.05670*, 2020.

Li, Jun, Antonio Moreno, and Dennis Zhang, "Pros vs joes: Agent pricing behavior in the sharing economy," *Ross School of Business Paper*, 2016.

Liu, Meng, Erik Brynjolfsson, and Jason Dowlatabadi, "Do digital platforms reduce moral hazard? The case of Uber and taxis," Technical Report, National Bureau of Economic Research 2018.

Ma, Hongyao, Fei Fang, and David C Parkes, "Spatio-temporal pricing for ridesharing platforms," *arXiv preprint arXiv:1801.04015*, 2018.

Sundararajan, Arun, *The sharing economy: The end of employment and the rise of crowd-based capitalism*, MIT Press, 2016.

Taylor, Terry A, "On-demand service platforms," *Manufacturing & Service Operations Management*, 2018, *20* (4), 704–720.

Varian, Hal R., "Computer Mediated Transactions," *American Economic Review*, 2010, *100* (2), 1–10.

Yan, Chiwei, Helin Zhu, Nikita Korolko, and Dawn Woodard, "Dynamic pricing and matching in ride-hailing platforms," *Naval Research Logistics (NRL)*, 2020, *67* (8), 705–724.

Zoepf, Stephen M, Stella Chen, Paa Adu, and Gonzalo Pozo, "The economics of ride-hailing: Driver revenue, expenses and taxes," *CEEPR WP*, 2018, *5*.

A Additional details on the experiment and quasi-experiment

A.1 Tests for pre-randomization covariate balance

One way to verify that the randomization was performed correctly is to test for systematic differences in pre-treatment observables between the experimental and the treatment groups. Table 1 reports balance tests for the experiment, in the form of p-values for the two-sided t-tests of the null hypothesis of no difference in means across the experimental groups. The experimental groups are seemingly well-balanced across various pre-experiment covariates, from which we can conclude that the randomization was performed correctly. An alternative way to check that the randomization was performed correctly is can be found in Figure 7, which reports pre-experiment weekly averages of several outcome variables.

Table 1: Summary statistics and mean comparison for providers in the experimental groups for pre-randomization observable variables.

	T0	T1	T2	T0-T1	T0-T2	T1-T2
observable	mean	mean	mean	p-value	p-value	p-value
	(s.e.)	(se)	(se)			
age	0.07	-0.23	-0.14	0.7	0.76	0.93
	(0.31)	(0.7)	(0.6)			
tenure	0	0	0	0.96	1	0.97
	(0.02)	(0.06)	(0.06)			
availability	0.63	0.65	0.63	0.41	0.7	0.71
	(0.01)	(0.02)	(0.02)			
cancelation rate	0.21	0.18	0.2	0.22	0.81	0.45
	(0.01)	(0.02)	(0.02)			
revenue	1.4	-8.69	2.96	0.77	0.97	0.82
	(17.09)	(30.8)	(40.49)			
hours rented	104.11	103.42	100.55	0.93	0.62	0.75
	(3.29)	(6.54)	(6.31)			

Notes: This table reports means and standard errors of various provider attributes across the three experimental groups, at the time that providers were allocated to treatment groups. These attributes are (i) age, (ii) tenure on the platform, (iii) vehicle availability, (iv) transaction cancelation rate, (v) revenue, and (vi) hours rented. Attributes (i), (ii) and (v), and (vi) are demeaned for confidentiality purposes. Attributes (iii) to (vi) are defined in Section 3, and are measured using observations for the month prior to the commencement of the experiment; other sample periods yield similar results. The reported p-values are for the two-sided t-tests of the null hypothesis of no difference in means across groups. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : * * *$.

A.2 Pricing interfaces

Figure 6 shows the pricing interfaces for the two treatment groups. The centrally set price is shown through a graph that depicts the hourly rate fluctuations. The platform utilized tooltips that provide more information, in order to minimize provider confusion over the change. The price slider and the associated tool tip, explaining how the slider works, is the only difference of this version with the fully centralized variant of the feature.

Figure 6: Pricing interfaces for the two treatment groups.

(i) Treatment group T1.

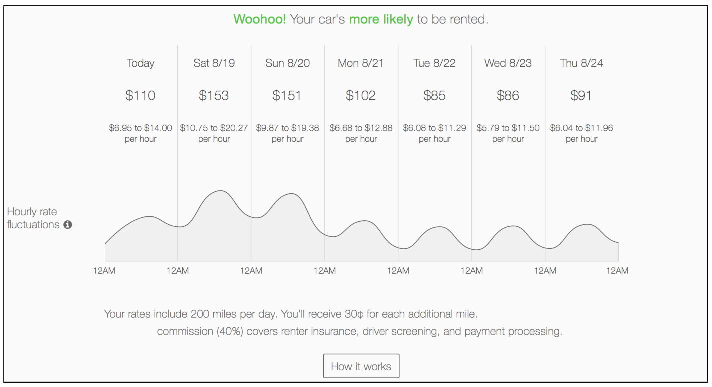

(ii) Treatment group T2.

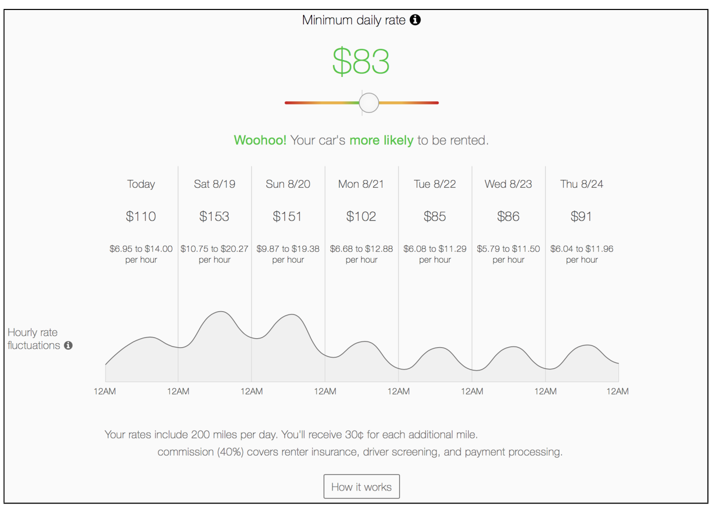

Notes: This figure shows the pricing interface for providers in the two treatment groups. The curves show the centrally set hourly prices for the next 7 days, and the minimum and maximum daily price and the daily price are presented in text. Providers in the treatment group T2 have access to a price slider that they can use to increase or decrease the centrally set price by up to 30 percent.

A.3 Alternative representations of the experimental results

Section 3 provides point estimates for the average treatment effects using a cross section of our data during the experimental period. To further evaluate the validity of our results, we provide in what follows alternative representations of the main outcome variables.

Figure 7 provides a longitudinal representation of the main outcome variables for each experimental group, by plotting their weekly averages for each treatment group, for each week before and during the experiment. This representation allows us to see diverging patterns clearly for outcome variables for which the treatment had a substantial effect. It also allows us to verify that the experimental groups are well balanced, and that the observed effects were not due to substantial changes in the responses and outcomes of the untreated providers, but rather due to changes in the treated providers.

It is worth noting that Figure 7 partially rules out provider communication from the set of potential threats to the experiment's internal validity: while it is impossible to rule out completely the possibility that subjects communicated with each other, we do not observe any significant shift in untreated providers' responses—availability, cancelations, and price changes. As such, we find no evidence that providers responded to the shift in competition by changing their economic behavior. This finding also supports the platform's claim that centralized pricing may increase transaction efficiency—in our data, providers seem to not be responsive to market changes, at least in the short-run.

Figure 8 provides a distributional representation of tte main outcome variables for each experimental group, by plotting the the kernel density estimate of each outcome variable for each treatment group. We can see substantial shifts in the outcome distributions, commensurate with the point estimates that are provided in the main body of the paper. For estimates employing alternative samples, seeAppendix B.

A.4 Subpopulation analyses

In addition to reporting the results of our analyses of the experiment for the entire population sample, we also report results for the following two subpopulations: (1) NEW, comprising providers with less-than-median tenure on the platform, and (2) LOW, comprising providers with less-than-median vehicle availability in the year prior to the experiment. We choose these subpopulations to understand how providers who are new to the platform, and users who are "casual"—users with low vehicle availability—potentially differ in their reactions compared to the overall population. Presumably, providers who are new or casual would react less strongly to the pricing mechanism change compared to long-time users and "serious" users.

The results of our analysis are presented in the regression tables of Appendix B. Providers belonging to different subpopulations exhibit very similar responses to the pricing change, suggesting that familiarity or inertia have little impact on the responses of providers.

Figure 7: Longitudinal representation of experimental effects on outcome variables.

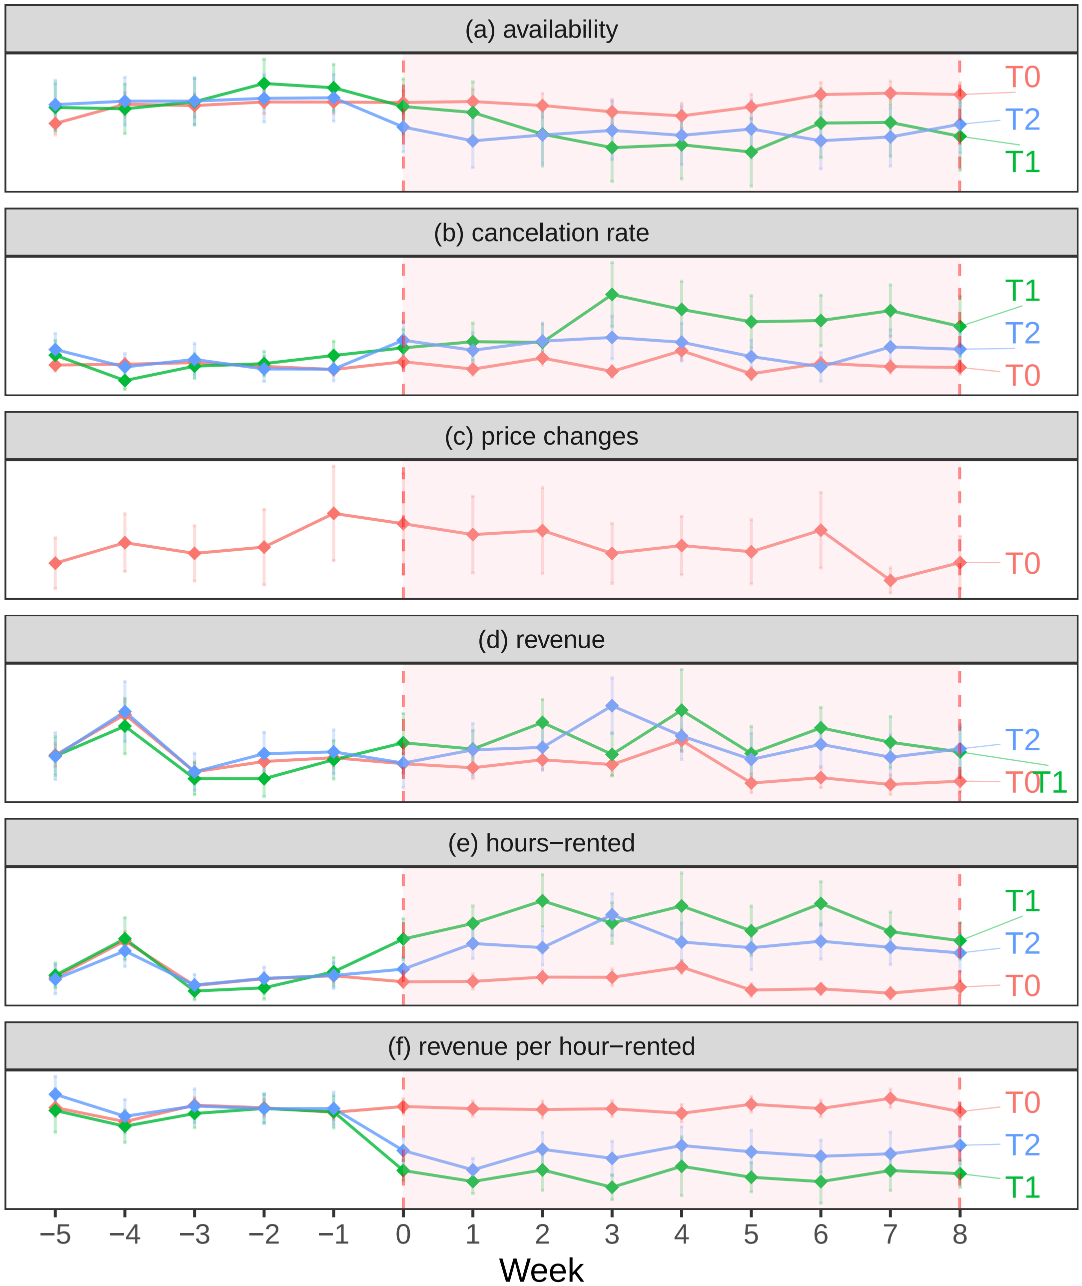

Notes: This figure reports a longitudinal representation of the effects of the treatment on providers' responses and economic outcomes. For each period, an estimate of the mean and a 95% confidence interval is plotted for each treatment, with the sample consisting of providers who were active during that period. The red vertical dashed lines and the pink-shaded area indicate the beginning and end of the experimental period. For point estimates, see Figure 2.

Figure 8: Distributional representation of experimental effects on outcome variables.

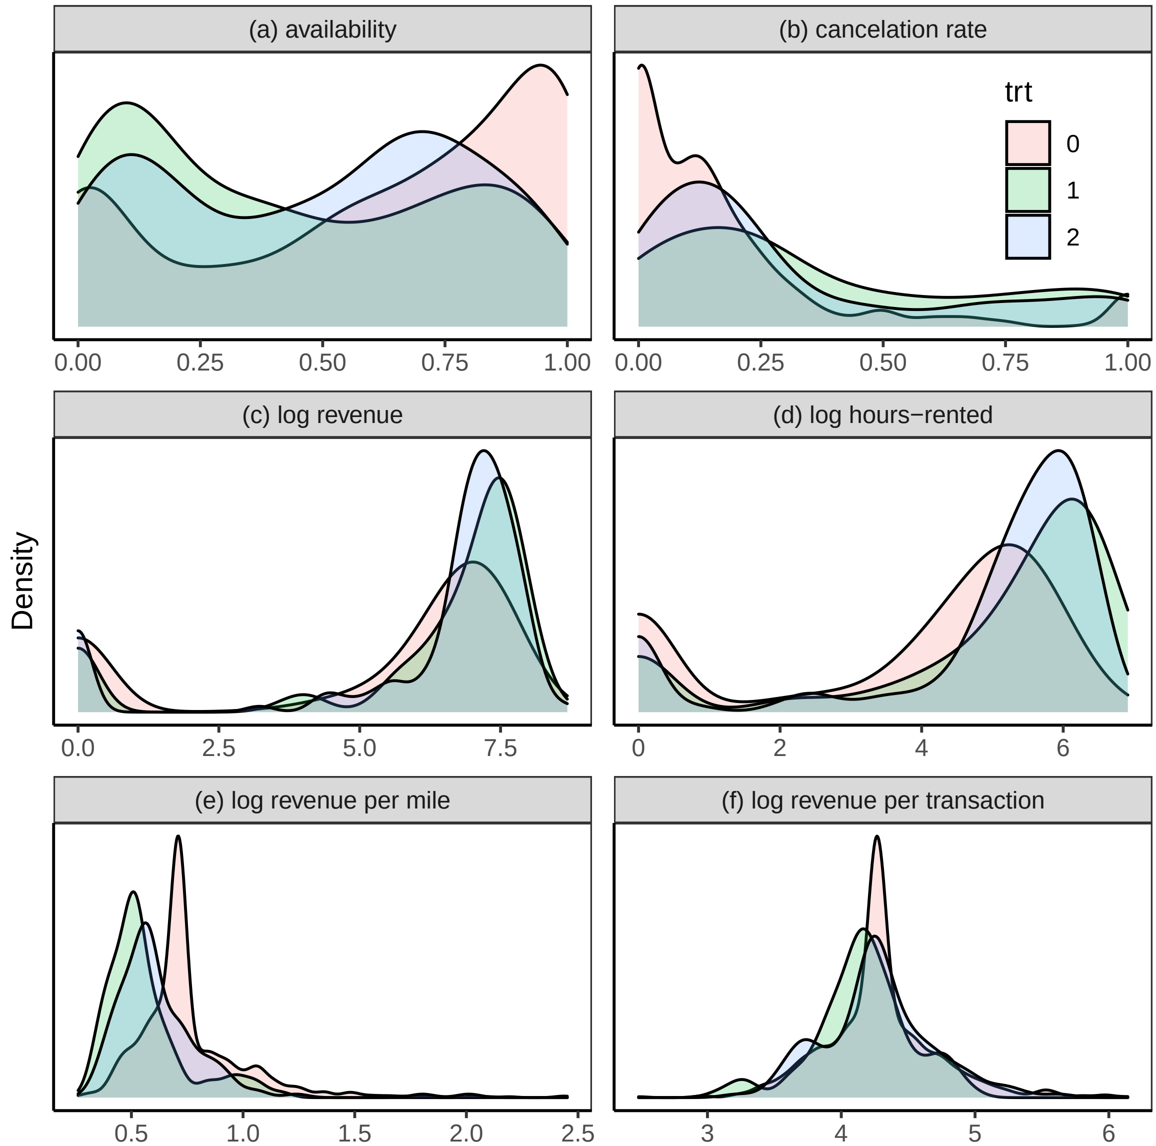

Notes: This figure reports a distributional representation of the effects of the treatment on providers' responses and economic outcomes. For each outcome and each experimental group, the kernel density estimate is plotted, with the sample constructed in the same manner as in Figure 2 and Figure 3, respectively. All outcome variables that cannot be naturally expressed as percentages are reported as logarithms. The dependent variables are calculated for the experimental period, and are (a) vehicle availability, (b) transaction cancelation rate, (c) provider revenue, (d) vehicle hours-rented, and (e) revenue per mile each car was driver, and (f) revenue per transaction

A.5 Pre-“treatment” trends

The credibility of the difference-in-differences approach of Section 5 depends on the suitability of the prior year as an appropriate counterfactual, and more specifically, on whether the time series of the outcomes in the two years exhibit parallel pre-treatment trends. Figure 4i in Section 5 provided visual evidence in support of the parallel trends assumption for aggregate outcomes—platform revenue and hours-rented. We present additional evidence in this section.

Figure 9 reports estimates of the mean outcome differences between the actual and the prior years for each week in the pre-treatment period. For each week-outcome pair, we report an estimate of the mean difference along with a 95% confidence interval, and we normalize all differences by subtracting the value of the first mean difference in each time series. Clearly, there is no deviation from a common trend in the two time series, and the results are consistent with the parallel trends assumption.

Figure 9: Evaluating the parallel trends assumption.

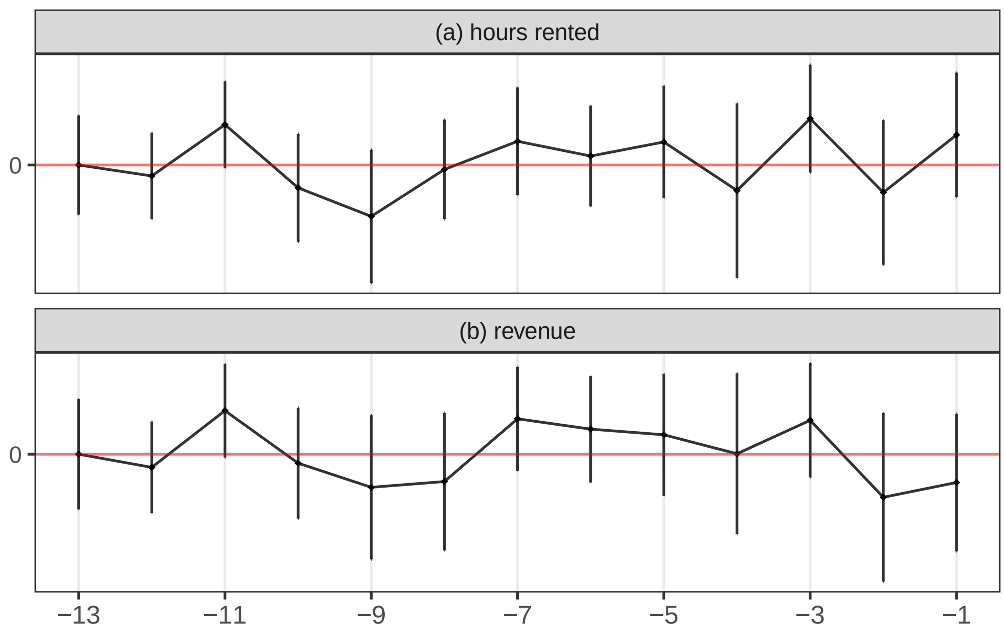

Notes: This figure evaluates the parallel trends assumption during the pre-treatment period for the year when the market-wide imposition took place and the previous year. Each panel plots an estimate of the mean difference between the dependent variable, along with a 95% confidence interval is plotted around each estimate. The dependent variables are calculated for each week during the pre-period, are normalize by the first value in each time series, and are (a) the hours rented, and (b) the revenue for each provider. More details on the design of the quasi-experiment are provided in Section 5.

Next, we follow Angrist and Krueger (1999) and conduct a direct pre-treatment test for parallel trends by estimating the regression

$$Y_{it} = \alpha_i + \delta t + \omega t \textsc{Trt}_i + \epsilon_{it}, \quad \text{for } t < 0, \tag{A1}$$

using data from all users during the pretreatment period across both years. In this regression, i is the user, t is the time period (week), the independent variable $\textsc{Trt}_i$ is equal to one during the actual year and zero otherwise, and Y_i is the outcome variable. As such, the parameter δ is the common trend, the parameter ω is the deviation from the common trend, and we are interested in testing against $H_0 : \omega = 0$, that is, against the null hypothesis of parallel trends. The results of this analysis are reported in Table 2. We fail to reject the null of parallel trends, and the estimated mean differences are small compared to the baseline outcomes. The results of this analysis are robust to employing alternative panel lengths, and to removing the user fixed effects. Taken together, our analyses confirm the validity of the difference-in-differences identification strategy.

Table 2: Evaluating the parallel trends assumption

	Dependent variable:	
	Revenue	Hours rented
	(1)	(2)
Trend (δ)	2.371***	0.417***
	(0.500)	(0.119)
Trend $\times$ Treatment (ω)	-0.282	0.149
	(0.668)	(0.158)
Constant	108.805***	24.679***
	(11.243)	(2.666)
R^2	0.004	0.008
Adjusted R^2	0.004	0.008

Notes: This table reports regressions testing the null of parallel trends in the pre-period of the quasi-experimental analysis. For more details, see the discussion in Appendix A.5. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

A.6 Post-experiment survey

The platform administered a short survey after the end of the experiment. The goal of this survey was to elicit additional feedback, in order to corroborate the findings of the experiment, as well as to better understand the reasons behind the providers' responses.

The survey was administered to providers that were assigned to the two treatment groups. These providers were asked (i) whether they preferred setting their own prices or the new pricing mechanism, (ii) whether they reacted to the centrally set prices by reducing the availability of their vehicles, (iii) and whether they reacted to the centrally set prices by canceling booked transactions. In addition to questions on a fixed scale, providers were also asked to provide free-form textual feedback describing their experiences with centralized new pricing mechanism. The response rate was identical across the two treatment groups, and equal to about 31 percent.

Figure 10 reports a bar plot for the responses of the subjects belonging to the two treatment groups. Each question is depicted in the caption of each panel, and the set of possible answers is depicted in the text of the horizontal axis. For the satisfaction question "Which of the following best describes your experience with the new pricing feature?" which is depicted in the top panel, the answers are ordered in increasing positivity of sentiment from left to right. For the availability and cancelation questions, which are depicted in the middle and bottom panels respectively, the answers are ordered in increasing agreement with the statement from left to right.

Respondents reacted fairly negatively to the new feature in aggregate: 78% providers in the experimental group T1 preferred setting prices for their vehicles themselves, rather than having the platform set prices. However, providers who were allowed to retain some pricing control reported substantially more positive attitudes towards the new pricing mechanism, with only 48% percent of these providers preferring setting their own prices. Furthermore, providers in the experimental group T1 were more likely to reduce the availability of their assets, and to cancel booked transactions, compared to providers in the experimental T2. Taken together, the findings from the post-experiment survey reaffirm the main empirical findings of the paper.

Figure 10: Answers to post-experiment survey questions

(i) "Which of the following best describes your experience with the new pricing feature?"

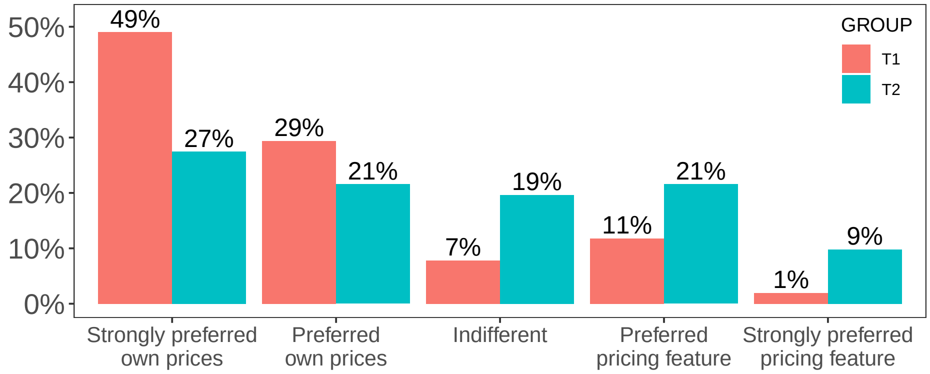

(ii) "When prices were set by the platform, I reduced the availability of my vehicles."

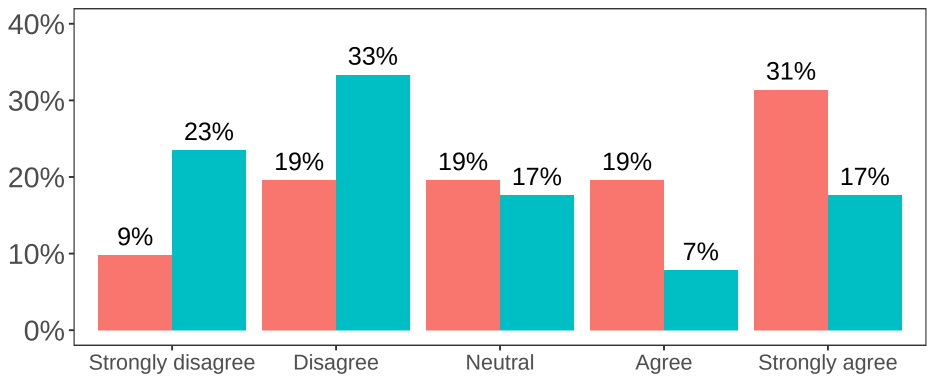

(iii) "When prices were set by the platform, I canceled some bookings after they were made."

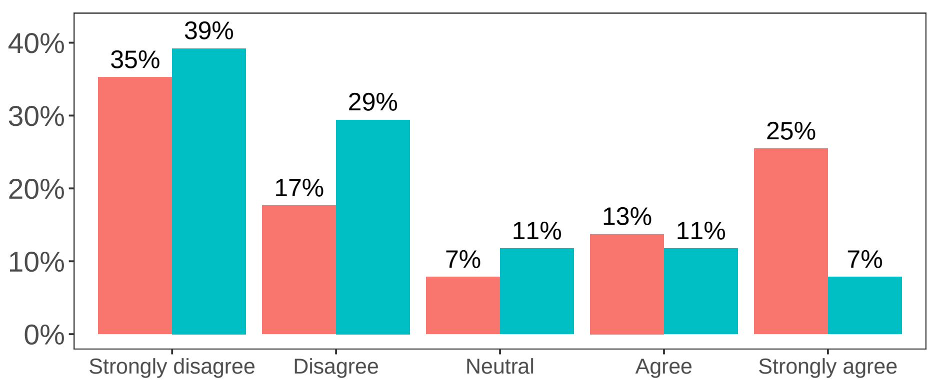

Notes: This figure plots the responses to questions administered to treated providers after the experimental period. The phrasing of each question is given in the caption of each panel. The possible answers to each question are given in the x-axis text of the respective panel.

B Regression tables

Table 3: Effects of the experimental treatment on provider exit.

	Dependent variable:		
	Provider Exit		
	ALL	NEW	LOW
	(1)	(2)	(3)
T1	0.305***	0.313***	0.225***
	(0.029)	(0.045)	(0.042)
T2	0.150***	0.088*	0.100*
	(0.029)	(0.042)	(0.043)
Constant	0.089***	0.119***	0.105***
	(0.011)	(0.017)	(0.016)
Observations	1,218	615	609
R^2	0.092	0.074	0.049

Notes: This table reports regressions where the dependent variable is provider exit from the platform. The independent variables are indicators for each experimental group, with the control group excluded. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 2 plots the treatment effects for the two treatment cells. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 4: Effects of the experimental treatment on availability.

	Dependent variable:					
	Vehicle availability					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	-0.071^{*}	-0.109^{*}	-0.029	-0.189^{***}	-0.235^{***}	-0.097^{*}
	(0.036)	(0.054)	(0.052)	(0.029)	(0.043)	(0.043)
T2	-0.070^{*}	-0.112^{*}	-0.038	-0.122^{***}	-0.138^{***}	-0.048
	(0.032)	(0.044)	(0.049)	(0.029)	(0.040)	(0.044)
Constant	0.645^{***}	0.658^{***}	0.472^{***}	0.612^{***}	0.615^{***}	0.445^{***}
	(0.012)	(0.017)	(0.018)	(0.012)	(0.016)	(0.017)
Observations	1,029	509	518	1,218	615	609
R^2	0.007	0.018	0.002	0.041	0.056	0.009

Notes: This table reports regressions where the dependent variable is vehicle availability. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 2 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (4)-(6). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 5: Effects of the experimental treatment on cancelation rate.

	Dependent variable:					
	Transaction cancellation rate					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	0.069**	0.079*	0.116**	0.198***	0.226***	0.196***
	(0.023)	(0.037)	(0.037)	(0.022)	(0.034)	(0.035)
T2	0.015	0.043	−0.012	0.126***	0.131***	0.094**
	(0.021)	(0.030)	(0.035)	(0.022)	(0.032)	(0.035)
Constant	0.165***	0.173***	0.170***	0.171***	0.181***	0.173***
	(0.008)	(0.011)	(0.013)	(0.009)	(0.013)	(0.014)
Observations	1,029	509	518	1,218	615	609
R^2	0.009	0.012	0.019	0.076	0.083	0.055

Notes: This table reports regressions where the dependent variable is cancelation rate. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 2 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (4)-(6). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 6: Effects of the experimental treatment on revenue.

	Dependent variable:					
	Provider revenue					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	446.671***	333.675*	399.000***	449.078***	386.514***	416.270***
	(87.339)	(133.334)	(102.752)	(68.898)	(101.956)	(83.807)
T2	431.492***	422.979***	442.289***	438.789***	420.315***	435.688***
	(78.653)	(107.170)	(97.518)	(68.549)	(95.181)	(85.604)
Constant	839.342***	850.849***	523.011***	784.269***	772.629***	486.589***
	(29.081)	(41.194)	(35.685)	(27.299)	(38.169)	(32.934)
Observations	1,029	509	518	1,218	615	609
R^2	0.047	0.038	0.058	0.056	0.046	0.068

Notes: This table reports regressions where the dependent variable is provider revenue. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 2 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 7: Effects of the experimental treatment on hours rented.

	Dependent variable:					
	Hours rented					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	188.697***	149.298***	157.855***	186.252***	164.019***	155.545***
	(17.839)	(26.635)	(21.735)	(14.101)	(20.453)	(17.734)
T2	139.847***	124.745***	132.511***	136.086***	120.658***	128.649***
	(16.064)	(21.409)	(20.628)	(14.029)	(19.094)	(18.114)
Constant	139.826***	142.295***	86.906***	130.445***	128.955***	80.657***
	(5.940)	(8.229)	(7.549)	(5.587)	(7.657)	(6.969)
Observations	1,029	509	518	1,218	615	609
R^2	0.140	0.104	0.140	0.161	0.130	0.155

Notes: This table reports regressions where the dependent variable is the number of hours providers' cars were rented. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 2 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 8: Effects of the experimental treatment on revenue per hour rented.

	Dependent variable:					
	Revenue per hour rented					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	-1.643^{**}	-1.655^{*}	-0.700	-1.543^{***}	-1.551^{**}	-0.507
	(0.502)	(0.772)	(0.682)	(0.392)	(0.587)	(0.553)
T2	$-0.808^{\ddagger}$	-0.753	0.488	-0.601	-0.568	0.551
	(0.452)	(0.620)	(0.648)	(0.390)	(0.548)	(0.565)
Constant	6.275^{***}	6.422^{***}	4.669^{***}	6.002^{***}	6.031^{***}	4.421^{***}
	(0.167)	(0.238)	(0.237)	(0.155)	(0.220)	(0.217)
Observations	1,029	509	518	1,218	615	609
R^2	0.012	0.011	0.004	0.013	0.012	0.003

Notes: This table reports regressions where the dependent variable is the revenue per hour a vehicle is rented. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 2 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3) Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 9: Effects of the experimental treatment on miles driven.

	Dependent variable:					
	Miles					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	140.481***	96.150***	123.170***	139.192***	118.233***	134.538***
	(13.673)	(22.105)	(16.572)	(13.677)	(21.056)	(19.183)
T2	102.925***	100.189***	87.016***	111.052***	99.447***	110.080***
	(12.313)	(17.767)	(15.728)	(13.608)	(19.657)	(19.594)
Constant	209.545***	211.570***	201.411***	153.900***	157.561***	110.609***
	(4.553)	(6.829)	(5.755)	(5.419)	(7.883)	(7.538)
Observations	1,029	509	518	1,218	615	609
R^2	0.132	0.083	0.130	0.108	0.075	0.104

Notes: This table reports regressions where the dependent variable is the miles a vehicle is driven. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 3 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 10: Effects of the experimental treatment on the number of transactions.

	Dependent variable:					
	Miles					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	0.796***	0.708***	0.469***	1.045***	1.003***	1.002***
	(0.113)	(0.191)	(0.122)	(0.129)	(0.198)	(0.172)
T2	0.495***	0.350*	0.390***	0.812***	0.562**	1.012***
	(0.102)	(0.153)	(0.116)	(0.128)	(0.185)	(0.176)
Constant	2.422***	2.473***	2.277***	1.724***	1.759***	1.161***
	(0.038)	(0.059)	(0.042)	(0.051)	(0.074)	(0.068)
Observations	1,029	509	518	1,218	615	609
R^2	0.060	0.033	0.043	0.070	0.048	0.087

Notes: This table reports regressions where the dependent variable is the number of transactions. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 3 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 11: Effects of the experimental treatment on revenue per mile.

	Dependent variable:					
	Revenue per mile					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	-0.442^{***}	-0.403^{***}	-0.366^{***}	-0.303^{***}	-0.274^{*}	-0.091
	(0.083)	(0.117)	(0.096)	(0.077)	(0.111)	(0.101)
T2	-0.286^{***}	-0.272^{**}	$-0.164^{\ddagger}$	-0.069	-0.041	0.133
	(0.075)	(0.094)	(0.091)	(0.077)	(0.104)	(0.103)
Constant	1.225^{***}	1.218^{***}	1.162^{***}	0.911^{***}	0.909^{***}	0.653^{***}
	(0.028)	(0.036)	(0.033)	(0.031)	(0.042)	(0.040)
Observations	1,029	509	518	1,218	615	609
R^2	0.036	0.034	0.030	0.013	0.010	0.005

Notes: This table reports regressions where the dependent variable is the revenue per mile a vehicle is used. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 3 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 12: Effects of the experimental treatment on revenue per transaction.

	Dependent variable:					
	Revenue per transaction					
	ALL (c)	NEW (c)	LOW (c)	ALL	NEW	LOW
	(1)	(2)	(3)	(4)	(5)	(6)
T1	-11.883^{**}	-15.877^{*}	$-11.817^{\ddagger}$	-0.818	-2.999	6.455
	(4.504)	(6.993)	(6.037)	(4.600)	(6.985)	(6.963)
T2	-2.007	0.517	-5.231	9.599^{*}	7.784	$12.314^{\ddagger}$
	(4.056)	(5.621)	(5.730)	(4.577)	(6.521)	(7.113)
Constant	79.473^{***}	78.693^{***}	79.929^{***}	58.404^{***}	58.370^{***}	45.214^{***}
	(1.500)	(2.161)	(2.097)	(1.823)	(2.615)	(2.736)
Observations	1,029	509	518	1,218	615	609
R^2	0.007	0.010	0.008	0.004	0.003	0.006

Notes: This table reports regressions where the dependent variable is the revenue per transaction made. The independent variables are indicators for each experimental group, with the control group excluded. Columns (1)-(3) compute the effects for providers who did not exit the platform. Columns (4)-(6) compute the effects for all providers. The subpopulation samples are defined as follows: ALL includes all providers, NEW includes providers with less-than-median tenure on the platform by the time that the experiment started, and LOW includes providers with lower-than-median vehicle availability. Figure 3 plots the treatment effects for the two treatment cells, for the estimation strategy of Columns (1)-(3). Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 13: Effects of price centralization on provider outcomes (window size = two weeks)

	Dependent variable:	
	Revenue	Hours rented
	(1)	(2)
Post	-113.027^{***}	-1.789
	(16.364)	(3.992)
Trt	-52.228^{***}	-9.689^{*}
	(15.696)	(3.858)
Post × Trt	$41.687^{\ddagger}$	13.648^{*}
	(21.947)	(5.368)
Constant	357.017^{***}	66.173^{***}
	(11.625)	(2.857)
R^2	0.012	0.002
Adjusted R^2	0.011	0.001

Notes: This table reports regressions where the dependent variable is provider revenue and hours rented, and the independent variables indicate whether the outcome is measured after the market-wide imposition of centralized pricing (Post), indicate whether the outcome took place during the year when the experiment took place (Trt), and an interaction of the two indicators (Post × Trt). The sample of observations is constructed by measuring dependent variable aggregates using a window size of two weeks. Figure 4 plots the treatment effect estimate as the percentage change over the difference-in-differences prediction. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 14: Effects of price centralization on provider outcomes (window size = four weeks)

	Dependent variable:	
	Revenue	Hours rented
	(1)	(2)
POST	-175.027^{***}	-22.217^{***}
	(24.737)	(5.589)
TRT	-70.320^{**}	-19.450^{***}
	(23.984)	(5.448)
POST × TRT	68.631^{*}	29.954^{***}
	(33.572)	(7.598)
Constant	608.697^{***}	117.823^{***}
	(17.549)	(3.986)
R^2	0.011	0.003
Adjusted R^2	0.010	0.002

Notes: This table reports regressions where the dependent variable is provider revenue and hours rented, and the independent variables indicate whether the outcome is measured after the market-wide imposition of centralized pricing (POST), indicate whether the outcome took place during the year when the experiment took place (TRT), and an interaction of the two indicators (POST × TRT). The sample of observations is constructed by measuring dependent variable aggregates using a window size of four weeks. Figure 4 plots the treatment effect estimate as the percentage change over the difference-in-differences prediction. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 15: Effects of price centralization on provider outcomes (window size = six weeks)

	Dependent variable:	
	Revenue	Hours rented
	(1)	(2)
Post	-297.069^{***}	-58.833^{***}
	(32.538)	(7.195)
Trt	-80.637^{*}	-30.189^{***}
	(32.022)	(7.109)
Post × Trt	114.508^{*}	51.318^{***}
	(44.626)	(9.879)
Constant	889.366^{***}	180.990^{***}
	(23.296)	(5.172)
R^2	0.016	0.009
Adjusted R^2	0.016	0.009

Notes: This table reports regressions where the dependent variable is provider revenue and hours rented, and the independent variables indicate whether the outcome is measured after the market-wide imposition of centralized pricing (Post), indicate whether the outcome took place during the year when the experiment took place (Trt), and an interaction of the two indicators (Post × Trt). The sample of observations is constructed by measuring dependent variable aggregates using a window size of six weeks. Figure 4 plots the treatment effect estimate as the percentage change over the difference-in-differences prediction. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 16: Effects of price centralization on provider outcomes (window size = eight weeks)

	Dependent variable:	
	Revenue	Hours rented
	(1)	(2)
POST	-303.865^{***}	-66.295^{***}
	(39.990)	(8.568)
TRT	$-67.854^{\ddagger}$	-35.737^{***}
	(39.472)	(8.489)
POST × TRT	134.798^{*}	65.984^{***}
	(55.033)	(11.803)
Constant	$1{,}093.926^{***}$	224.874^{***}
	(28.597)	(6.151)
R^2	0.010	0.008
Adjusted R^2	0.010	0.007

Notes: This table reports regressions where the dependent variable is provider revenue and hours rented, and the independent variables indicate whether the outcome is measured after the market-wide imposition of centralized pricing (POST), indicate whether the outcome took place during the year when the experiment took place (TRT), and an interaction of the two indicators (POST × TRT). The sample of observations is constructed by measuring dependent variable aggregates using a window size of eight weeks. Figure 4 plots the treatment effect estimate as the percentage change over the difference-in-differences prediction. Significance indicators: $p \leq 0.1 : \ddagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$, and $p \leq .001 : ***$.