Monopsony Power in the Gig Economy

Jack Fisher *

December 12, 2024

Abstract

Many workers provide services for customers via digital platforms that may exert monopsony power. Typical expositions of this phenomenon are inapplicable because platforms post prices to both sides of a two-sided market, and platform-specific labor supply is hard to measure when workers multi-app. This paper develops a model of a typical gig labor market that deals with these issues. Platforms exploit monopsony power to markup their commission rate and reduce equilibrium wages. A worker union sets the first-best commission rate when the customer market is competitive. I estimate the model using public data, including causal estimates from the literature on Uber's US ridesharing marketplace. The results imply the platform exploits labor market power to depress drivers' earnings but faces competition for passengers. An optimally set commission cap raises wages by 14 percent, but minimum wages on utilized hours harm workers.

*University of Virginia: jack.w.fisher.91@gmail.com. I am thankful to Johsua Angrist, Florian Ederer, Chiara Farronato, John Horton, Myrto Kalouptsidi, Matthew Leisten, Alex MacKay, Marc Rysman, Chris Stanton, Michael Sullivan, and Catherine Thomas for their helpful comments.

1 Introduction

Digital platforms, like Amazon's marketplace and Uber's ridesharing platform, often operate two-sided markets that facilitate exchanges between buyers and sellers (Rysman, 2009). In this setting, concerns over market power may arise because strong network effects mean participants benefit from being in the same marketplace, and individuals are atomistic relative to platforms. However, the very existence of network externalities can make platforms reluctant to exploit market power over one side of the market if it harms the other side. Consequently, it is unclear when these worries are warranted (Jullien et al., 2021).

This issue is prescient for labor economists studying monopsony power given the rise of gig work, where workers provide short-term and independent labor services via digital intermediaries (Dube et al., 2020).[1] In the US, a large and growing share of workers participate in the gig economy, and globally, hundreds of millions earn income via these types of arrangements (Anderson et al., 2021; Datta et al., 2023). Yet, policymakers lack a framework to understand how monopsony power manifests in the gig economy, where platforms post prices and commission rates rather than wages, and what institutional remedies may be effective.

This paper develops a tractable model to study a typical and important gig labor market: ridesharing (Garin et al., 2023). The model provides clear insights into platform pricing, the merits of alternative market designs, and their interplay with market power from both a platform and a social planner's perspective. Further, a small number of statistics are sufficient to estimate the model's parameters. I demonstrate the framework's utility by using it to evaluate the extent of market power enjoyed by the US's largest ridesharing platform, Uber, with publicly available data, including causal estimates from the academic literature that use proprietary microdata.

The empirical results indicate that Uber wields substantial monopsony power over drivers but faces strong competition for riders. The platform finds it profitable to charge drivers a high commission rate, which reduces their labor supply and increases waiting times. Passengers are not compensated for the latter with lower prices, so both drivers and riders would be better off absent monopsony power. Quantitatively, the commission rate is 15 percentage points higher than the social planner's optimum. If this were restored to its first-best level, wages would increase by 14 percent

[1] An agent has monopsony power when they can pay a lower unit price for a good or service if they buy a lower quantity. In this context, there is an ambiguity about whether platforms buy labor that they then offer customers or whether customers are the buyers and platforms only mediate exchanges. In what follows, this is only an issue of semantics.

after accounting for the platform's passenger pricing response. The theory reveals that delegating commission rate setting to an organization that aims to maximize wages, like a drivers' union, would deliver this outcome.

Concretely, I consider a ridesharing platform that operates a two-sided marketplace, and that sets the price of exchanges and the commission rate they receive. Riders on the platform care about the price they face and the utilization of drivers, which determines waiting times under a variety of micro-foundations. Hourly wages, which also depend on utilization since drivers are only paid when they have passengers in their vehicle, determine the supply of drivers.

The market reaches equilibrium through adjustments in worker utilization (Hall et al., 2023). After the platform's pricing is set, drivers enter and exit as their wage rate moves with utilization, and riders change their demand as waiting times fluctuate accordingly. Intuitively, this leads to a fixed point equilibrium condition that constrains the platform's profit maximization problem. That is, equilibrium utilization must equal the optimally chosen ratio of supply and demand, which itself depends on utilization.[2]

The model delivers clear intuitions about decision-making by intermediaries and how they contrast with a social planner. Three behavioral elasticities describe how platforms incorporate market power into their optimal choice of price and commission rate. First, the elasticity of driver supply to hourly earnings, which corresponds to the extent of monopsony power. Second, the elasticity of passenger demand to price, and third, the elasticity of passenger demand to utilization or waiting times. The latter two elasticities jointly reflect a platform's monopoly power, but their relative magnitudes are important for pricing.

Platforms markup the commission rate that drivers pay when they enjoy monopsony power. Equivalently, the platform reduces drivers' keep rate identically to markdowns in textbook wage-posting models (Manning, 2011). To this extent, the platform-specific labor supply elasticity is still a useful measure of monopsony power in this context, but it is an incomplete picture. Demand elasticities for price and waiting times, and a precise counterfactual, which dictates the response of rider prices and driver utilization, are necessary to infer the equilibrium impact on wages and welfare

[2]The model has implications outside of ridesharing for two reasons. First, the role of utilization is an alternative specification of network effects that is relevant to other markets. For example, product differentiation in consumer goods when prices are set centrally. In this case, the higher the price, the more products are supplied, and the more likely a consumer will find something close to their ideal variety. I thank Catherine Thomas for this clarifying comment. Second, in terms of empirical content, the strength of monopsony power is likely to be similar across different gig labor markets within the same geography if market power comes from a common source, such as unattractive options outside of the gig economy.

(Kroft et al., 2020; Van Reenen, 2024).

The relative magnitude of price and waiting time elasticities also determines commission rates; if waiting times are important to customers, the platform reduces its commission to encourage driver supply. This can be seen as a nuanced type of "see-saw" principle (Rochet and Tirole, 2003, 2006).[3] Conversely, the driver supply elasticity does not affect passenger prices, so the platform does not pass on the benefits of monopsony power to riders. The platform increases its commission rate instead of passing on savings to customers because price reductions trigger longer waiting times and dampen demand, which makes this strategy unattractive. Intuitively, the platform separately exploits monopoly and monopsony power with its prices and commission rates, respectively.

The platform's optimal pricing strategy and commission rate coincide with the social planner's when the platform has no market power. To the extent that occurs under platform competition, this result contrasts with ambiguous findings in the literature (Hagiu and Jullien, 2014; Tan and Zhou, 2021). This is because the ratio of participants on either side of the market (*i.e.*, utilization) mediates network effects (*i.e.*, waiting times).[4] Therefore, fragmenting participants across different platforms is inconsequential if the utilization rate is preserved, which can be interpreted as a multi-apping assumption that is readily applicable in ridesharing.

Setting the commission rate to maximize wages yields the first-best commission rate when the platform retains control of rider prices and faces competition for riders. From a policy perspective, this can rationalize commission caps (Sullivan, 2022) and delegating commission rate setting to a drivers' union with a mandate to maximize hourly earnings. This result is the two-sided market analog of the classic union remedy to monopsony power in one-sided labor markets (Robinson, 1969), and also provides a feasible counterfactual with which to compare outcomes in the *status quo* equilibrium.

The model yields transparent necessary conditions on demand elasticities for equilibrium existence, which have substantive implications for rich structural models of ridesharing markets (Castillo, 2023; Rosaia, 2020). In particular, the elasticity of rider demand with respect to price must exceed one, and this lower bound rises as passengers' value of time increases. When the condition does not hold, the platform always

[3] That is, the tendency for factors that depress prices on one side of the market to raise prices on the other side of the market.

[4] Typically, network effects are specified to depend solely on the number of participants on the other side of the market, which does not account for the possibility of network congestion (Weyl, 2010).

wants to increase prices because waiting times fall and offset the decline in demand too much for the platform to have an interior solution. Problematically, elasticity estimates often used to calibrate demand in ridesharing markets do not satisfy this condition (Cohen et al., 2016). This stems from the use of short-term price variation, which underestimates rider price sensitivity. However, reinterpreting long-run pricing experiments in Hall et al. (2023) through the model provides direct evidence that Uber's US marketplace does satisfy the conditions for equilibrium existence.

The small number of parameters makes the model easy to estimate with little information, which is especially valuable in a context where data is often proprietary. I illustrate this point with an application to the US's largest ridesharing platform, Uber. Evaluating the extent of Uber's market power over the rideshare industry is an important issue in its own right. In the US, there are 1.5 million drivers actively working on the platform and over six million globally. Moreover, despite evidence that many workers benefit substantially from the opportunity to partake in ridesharing markets and alike (Fisher, 2022), concerns remain about the welfare of individuals subject to these work arrangements (Prassl, 2018; Ravenelle, 2019).

One reason for this anxiety is a fear that platforms use monopsony power to benefit at the expense of drivers. Fears around monopsony power in the gig economy are particularly pronounced because of prevalent underemployment (Lachowska et al., 2023), which creates a captive pool of workers. Gig work also entails a unique bundle of amenities that other jobs cannot offer (*e.g.*, flexible hours, see Chen et al. (2019)), and many labor market regulations designed to ameliorate an imbalance of bargaining power between employers and employees do not apply in the gig economy.

Therefore, evaluating platform monopsony power over workers in ridesharing markets and potential remedies is a first-order policy question. Unfortunately, aside from two-sidedness, estimating monopsony power in this context is challenging for at least two reasons. First, microdata on the ridesharing industry is proprietary and researchers can only access this data at the behest of platforms.[5] To overcome this hurdle, I use the framework outlined above to reduce the data demands of estimation. Second, drivers' labor supply is hard to measure (Harris and Krueger, 2015; Hyman et al., 2020). Intermediaries typically record "online" hours, but this measure does not account for multi-apping across platforms or whether workers genuinely plan to supply their labor. I circumvent this issue by inferring a platform-specific labor supply elas-

[5]There are notable efforts to scrape data on ridesharing and related markets, for example, see Rosaia (2020). Efforts like the EU's Data Act may also be one way to avoid this issue.

ticity from pricing decisions. This does not require platforms to observe a measure of genuine labor supply, instead they can optimize through experimentation.

Uber is well suited to the analysis set out in this paper. The firm sets the price of rides that drivers complete for passengers on its platform and receives a share of the fare. Although the exact commission rate varies from ride to ride, in the words of Uber's CEO, Dara Khosrowshahi, "[the platform] optimizes for an average take-rate".[6,7] Moreover, passengers care about prices and utilization because of wait times, and naturally, earnings determine drivers' labor supply. Thus, the platform's pricing in the marketplace and its participants satisfy the model's core assumptions.

To identify the driver and rider elasticities that Uber faces, I use public information on the platform's commission rate, costs, prices, and the elasticity of equilibrium utilization to price. These numbers apply to the US rideshare marketplace around 2017. The model definition of a commission rate is the fraction of the fare riders pay that drivers do not receive. This differs from the 20 or 25 percent fee that Uber publicizes because some of the fare, known as the booking fee, is earmarked to cover costs. Using proprietary data reported in Castillo (2023) and Cook et al. (2021), I calculate a commission rate of 34 percent.[8] The same sources imply a marginal cost of mediation equal to 18 percent of the fare, which is made up of transaction fees, sales tax, and insurance premiums. To account for uncertainty, I consider a range of both commission rates and costs.

The final moment to identify the model is the elasticity of equilibrium utilization to price, which Hall et al. (2023) estimates using pricing experiments conducted by Uber between 2014 and 2017 in US cities. In response to a 10% price increase, utilization eventually falls by 14%. This response is measured up to six months after the treatment was introduced and is informative of the *long-run* behavioral elasticities that drive base pricing decisions. Further, the magnitude of the response verifies the aforementioned necessary condition on demand elasticities holds in Uber's marketplace since price rises reduce demand more than one-for-one in percentage terms. This moment is reported with a standard error that I incorporate into the estimation.

Bringing the model to the data suggests that Uber exerts substantial monopsony power over drivers but faces strong competition for riders. The central scenario im-

[6] See Dara's interview with The Rideshare Guy here.

[7] The "take rate" is another phrase for the commission rate.

[8] Notably, Uber's commission rate was increasing up to 2017 (Caldwell and Oehlsen, 2021), and public financial filings indicate the commission rate may have risen since. The ratio of revenue to gross bookings in Uber's 10-K filings increased by around one-third in 2022, from approximately 20 percent to 27 percent, although accounting practices and the growth of Uber Eats make this hard to interpret.

plies that the platform faces a precisely estimated elasticity of driver supply to wages of 4.27.[9] Viewed through the model, this implies a 15 percentage point markup of the commission rate relative to the competitive benchmark which, in standard wage-posting models, corresponds to a wage markdown of one-fifth. However, in a two-sided market, prices and utilization also change in a counterfactual equilibrium. Therefore, it is necessary to consider a precise counterfactual to understand the impact of monopsony power on wages and worker welfare.

I consider setting commission rates to their first-best level noting that, given the competitive rider market, this is practically feasible by setting commission rates to maximize wages and leaving Uber in charge of rider prices. In this scenario, a commission rate fall of 15 percentage points triggers the platform to raise prices by almost half. In turn, this causes utilization to fall by two-thirds so that, overall, wages rise by 14 percent. This prediction differs significantly from a one-sided model of monopsony power and demonstrates the importance of accounting for two-sidedness. Under an iso-elasticity assumption, these wage gains precipitate a 20 percent increase in worker welfare after adjusting for gig work's share of total income.

Lastly, I study a minimum wage applying to utilized hours (*i.e.*, passenger-in-car hours) for two reasons. First, Uber released public information on its blog describing the equilibrium impact of such a policy in Seattle. This serves as an out-of-sample test for the model, which it passes. Second, this policy has been popular amongst state and local governments (*e.g.* see recent legislation in Minnesota) but lacks a formal motivation. The model shows these policies are appealing when the platform's markup on rider prices is large. A high markup indicates a low price elasticity of demand so that when the platform raises prices in response to the policy, passengers absorb the higher cost of ridesharing and driver wages increase in equilibrium. This condition is not met in the data, which suggests these policy efforts have likely harmed workers.

This paper proceeds as follows: section 2 develops a model of a ridesharing market, section 3 considers alternative marketplace designs, section 4 presents the empirical application to Uber, section 5 assesses the impact of Uber's market power on workers' wages and welfare, and section 6 concludes.

[9] This number is close to estimates found across US labor markets in Lamadon et al. (2022).

2 Model of a Two-Sided Ridesharing Marketplace

This section develops a model of a two-sided ridesharing marketplace operated by a platform. The theory builds upon the framework of Hall et al. (2023) by explicitly considering the platform's price and commission rate setting.

2.1 Market Participants, Wages, and Equilibrium

This subsection describes the decisions of the different agents who interact in the marketplace and the definition of equilibrium in the model.

Drivers. Ridesharing drivers decide how much to work on the platform according to the wage rate that they can earn. An aggregate driver labor supply function $H(w)$, which depends on hourly wages w, determines the number of driver hours available to the platform. The function comprises extensive and intensive margin labor supply responses to changes in earnings. In ridesharing markets, intensive margin labor supply responses extend beyond choosing how many hours to work conditional on working. For example, intensive margin responses may include how devoted workers are to the platform, which can take the form of geographical positioning and acceptance rates. In this sense, $H(w)$ reflects workers' *platform-specific* labor supply.

Riders. Passengers demand hours of transportation which is described in aggregate via a demand function $D(p, x)$. Their demand depends on the price of an hours ridesharing services p and driver utilization x, which determines waiting times. This assumption has two alternative micro-foundations. First, under a constant returns-to-scale matching function between drivers and riders, waiting times are solely a function of the utilization rate and the matching technology's structural parameters (Cullen and Farronato, 2021). With access to proprietary data, Hall et al. (2023) argues that this is a reasonable approximation in the context of Uber. Second, queuing theory finds utilization is crucial in determining waiting times, most famously in Kingman's equation (Kingman, 1961). Here, the structural parameters that determine waiting times correspond to features of the distribution of arrivals and characteristics of trips.

Wages. Hourly wages are an equilibrium quantity. They depend on the price per hour of transportation p that the platform charges, the fraction of fares that drivers

retain θ (*i.e.*, the keep-rate or one minus the commission rate), and the proportion of supply hours that drivers are transporting passengers x (*i.e.*, the utilization rate). Taken together, hourly earnings are given by

$$w = p \cdot \theta \cdot x. \tag{1}$$

Equilibrium. The marketplace equilibrates through adjustments in utilization after the platform has set its prices; drivers enter and exit as their wage rate moves with utilization while riders also change their demand as waiting times fluctuate. In particular, given a price and a commission rate, equilibrium requires that

$$x = \frac{D(p,x)}{H(p \cdot \theta \cdot x)}. \tag{2}$$

In other words, utilization must satisfy a fixed point such that equilibrium utilization equals the ratio of optimally chosen demand and supply of ridesharing hours, which also depend on utilization. This is analogous to the condition in Hall et al. (2023). For a given p and θ, a unique equilibrium exists if $\frac{\partial H(w)}{\partial w} > 0$, $H(w) \geq 0$, $\frac{\partial D(p,x)}{\partial x} < 0$, and $D(p,x) \geq 0$, which I assume for the remainder of the paper.

2.2 The Platform

A platform selects a price and commission rate to maximize profits but is constrained by the equilibrium adjustment of driver utilization, which also encapsulates driver and rider behavior. Formally, platforms face the following problem

$$\max_{p,\theta} \ \big[p \cdot (1-\theta-\tau) - c\big] \cdot D(p,x) \ \text{ subject to } \ D(p,x) = x \cdot H(p \cdot \theta \cdot x), \tag{3}$$

where τ represents costs that are proportional to the fare (*e.g.*, taxes and transaction fees) and c denotes other costs of mediation (*e.g.*, insurance premiums). Platform optimization yields two first-order conditions (4) and (5) for p and θ, respectively,

$$1 + \mu^* \cdot \left(\varepsilon_{D,x} \cdot \varepsilon_{x,p} - \varepsilon_{D,p}\right) = 0, \tag{4}$$

$$-\frac{\theta^*}{1-\theta^*-\tau} + \mu^* \cdot \varepsilon_{D,x} \cdot \varepsilon_{x,\theta} = 0, \tag{5}$$

where $\varepsilon_{D,x} = -\frac{\partial D(\bullet)}{\partial x} \cdot \frac{x}{D(\bullet)}$, $\varepsilon_{D,p} = -\frac{\partial D(\bullet)}{\partial p} \cdot \frac{p}{D(\bullet)}$, $\varepsilon_{x,p} = -\frac{dx}{dp} \cdot \frac{p}{x}$, $\varepsilon_{x,\theta} = -\frac{dx}{d\theta} \cdot \frac{\theta}{x}$, and $\mu = \frac{p \cdot (1-\theta-\tau)-c}{p \cdot (1-\theta-\tau)}$.[10] The latter term is a Lerner-type index (Lerner, 1934), which equals

[10] For ease of interpretation, I sign all elasticities to ensure that they are positive.

the share of platform revenue that is profited from one hour of ridesharing after the platform pays drivers, taxes, and fees. Asterisks denote optimally chosen endogenous variables.

Equation (4) reveals that raising prices mechanically increases revenue but simultaneously impacts demand via two behavioral channels. First, higher prices reduce demand in the traditional sense. Second, higher prices raise wages, which encourages higher driver supply and, in turn, reduces utilization and increases demand. Equation (5) follows an analogous logic for the setting of commission rates. Raising the commission rate leads to more revenue but also increases utilization due to lower wages that discourage driver supply and, eventually, decrease demand.

Comparative statics on the equilibrium condition described by equation (2) provide two more equalities that connect the demand and supply elasticities

$$\varepsilon_{x,p} = \frac{\varepsilon_{D,p} + \varepsilon_{H,w}}{\varepsilon_{D,x} + 1 + \varepsilon_{H,w}}, \tag{6}$$

$$\varepsilon_{x,\theta} = \frac{\varepsilon_{H,w}}{\varepsilon_{D,x} + 1 + \varepsilon_{H,w}}, \tag{7}$$

where $\varepsilon_{H,w} = \frac{\partial H(\bullet)}{\partial w} \cdot \frac{w}{H(\bullet)}$. Equilibrium utilization responds more strongly to a change in price than to a change in the commission rate because the former affects both drivers and riders directly. Intuitively, the numerator of equations (6) and (7) reflect the direct effect of their respective price and commission rate changes, while the denominators capture equilibrium effects.

Theorem 1 (The platform's optimal pricing)**.** *The platform's optimal price and commission rate can be expressed as a function of elasticities that describe driver and passenger behavior as follows*

$$p^* = \frac{1}{1 - (\frac{1}{\varepsilon_{D,p}} + \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}})} \cdot \frac{c}{1-\tau}, \tag{8}$$

$$1 - \theta^* = 1 - (1-\tau) \cdot \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} \cdot \frac{\varepsilon_{H,w}}{1 + \varepsilon_{H,w}}, \tag{9}$$

when $\varepsilon_{H,w} > 0$, $\varepsilon_{D,x} > 0$, *and* $\varepsilon_{D,p} > 1 + \varepsilon_{D,x}$.

Proof. Substituting equations (6) and (7) into the first-order conditions (4) and (5), and rearranging gives expressions (8) and (9). □

Below, I treat $\varepsilon_{D,p}$, $\varepsilon_{D,x}$, and $\varepsilon_{H,w}$ as structural parameters that are invariant to counterfactual scenarios, which yields an unrestrictive approximation to equilibrium

outcomes. To better understand the implications of Theorem 1 and other results below, two formal definitions are helpful.

Definition 1 (Perfect competition for drivers)**.** *$\varepsilon_{H,w}$ converges to infinity.*

Definition 2 (Perfect competition for riders)**.** *Both $\varepsilon_{D,p}$ and $\varepsilon_{D,x}$ converge to infinity, and $\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}$ converges to κ. This implies that $\varepsilon_{D,p} - \varepsilon_{D,x}$ converges to infinity.*

The definition of perfect competition for drivers is straightforward; the platform has no monopsony power when drivers are infinitely sensitive to changes in their wage. For riders, perfect competition implies that they are infinitely sensitive to changes in the price and waiting times. But this does not define the ratio of or difference between these elasticities. To resolve this, I assume that the platform's Lerner index converges to zero under perfect competition on both sides of the market, which requires $\varepsilon_{D,p} - \varepsilon_{D,x}$ and $\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}$ converge to infinity and a constant, respectively.

Price. The platform determines the price for passengers by marking up marginal costs according to the rider-side behavioral elasticities. Two factors affect the markup. First, if riders are elastic to waiting times relative to price, which is captured by the $\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}$ term in equation (8), then the platform sets a higher price. This is socially efficient; the platform is internalizing the congestion effects of lower prices on the network. Second, if the platform has some monopoly power then it sets higher prices, and this is socially inefficient.

Pricing equation (8) is noticeable for not including the platform-specific labor supply elasticity. This implies that the platform does not use monopsony power to benefit passengers in terms of lower prices. In other words, there is no "seesaw" effect for rider prices (Rochet and Tirole, 2003, 2006). The platform prefers to increase its commission rate instead of passing on savings to customers because price reductions trigger longer waiting times and dampen demand, which makes this strategy unattractive. In other words, conditional on an optimal commission rate, the price is set to maximize the total revenue extracted from riders. This has important anti-trust implications because abusing labor market power damages both sides of the market, which is the relevant legal test set by the US Supreme Court in two-sided marketplaces.[11]

Commission rate. The clearest implication of equation (9) is that platforms use monopsony power to raise their commission rate through the term $\frac{\varepsilon_{H,w}}{1+\varepsilon_{H,w}}$. All else

[11] Ohio v. American Express Co., 585 U. S. 529, (2018).

equal, this is equivalent to reducing workers' wages by $\frac{1}{1+\varepsilon_{H,w}}$ percent—the same markdown as in one-sided labor market models of monopsony power with wage-posting (Manning, 2011). However, in two-sided markets, commission rate markups do not directly translate to wage markdowns because there are pricing responses on the other side of the market and equilibrium effects on utilization. I explore these mechanisms in section 3.

Interestingly, commission rates will not necessarily only recoup marginal costs absent monopsony power. In particular, the commission rate under perfect competition for drivers equals

$$\lim_{\varepsilon_{H,w}\to\infty} 1-\theta^* = 1-(1-\tau)\cdot\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}. \tag{10}$$

In this instance, the ratio of demand elasticities also determines the commission rate. If rider demand is more sensitive to waiting times than price, then commission rates are kept high to incentivize drivers to provide capacity on the platform. The platform can still charge commission without any monopsony power because it must recoup costs and wages are not monotonically increasing in the commission rate. In subsection 3.1, I show that the commission rate implied by equation (10) maximizes the wage rate when there is perfect competition for riders.

Markups. The model also yields an augmented Lerner rule. If the platform faces a perfectly competitive market for drivers, then the optimal markup is given by

$$\lim_{\varepsilon_{H,w}\to\infty} \mu^* = \frac{1}{\varepsilon_{D,p}-\varepsilon_{D,x}}. \tag{11}$$

This combines the traditional pricing motivation of a monopolistic firm in a one-sided environment with an additional two-sided market concern. That is, increasing prices reduces utilization which partially offsets the fall in demand and, therefore, justifies a higher markup from a profit-maximizing perspective.

Behavioral elasticities. Looking at equation (6), the conditions on demand elasticities are equivalent to requiring that $\varepsilon_{x,p}$ exceeds one. This is directly tested by Hall et al. (2023) (see their Figure 5) and holds comfortably in the long run. If this condition does not hold, the platform does not have an interior solution to its problem and always wants to raise prices. For intuition, consider the condition not holding and keeping the commission rate fixed. An increase in the price necessarily increases

drivers' wages since the relative decline in utilization is less than the rise in the price. If wages are higher, then profits are higher too, once utilization has not declined too much—which failure of the condition ensures.

These conditions have substantive implications for rich structural models of ridesharing. If long-run demand elasticities are calibrated to match short-run responses to price changes (*e.g.*, Cohen et al. (2016)), which do not satisfy Theorem 1, then it is not possible to model a profit-maximizing platform. To circumvent this issue, and motivated by the fact platforms are competing for centrality, the prior literature has assumed that drivers' and riders' welfare features in the platform's objective function (Castillo, 2023; Rosaia, 2020; Sullivan, 2022).

2.3 The Social Planner

The platform's pricing can differ from the social optimum because it may exert market power over either side of the market. Socially efficient pricing maximizes the sum of platform profits and rider and driver surplus subject to participants' incentives, which are embedded in the equilibrium condition. Formally, under an iso-elastic assumption, the social planner faces the following problem

$$\max_{p,\theta} \ \left[p \cdot (1-\theta-\tau) - c\right] \cdot D(p,x) + \frac{p \cdot D(p,x)}{\varepsilon_{D,p} - 1} + \frac{w \cdot H(w)}{1+\varepsilon_{H,w}}$$
$$\text{subject to } \ D(p,x) = x \cdot H(p \cdot \theta \cdot x). \tag{12}$$

That is the social planner places equal weight on the platform's profits, rider surplus, and consumer surplus. The social planner's objective function can be rewritten as an alternative parameterization of the platform's problem after incorporating the equilibrium constraint as follows

$$\left[p \cdot \left(\frac{\varepsilon_{D,p}}{\varepsilon_{D,p}-1} - \frac{\varepsilon_{H,w}}{1+\varepsilon_{H,w}} \cdot \theta - \tau\right) - c\right] \cdot D(p,x), \tag{13}$$

which leads to the following result.

Theorem 2 (Efficient competitive private equilibrium)**.** *The private equilibrium, which is described by equations* (8) *and* (9)*, is socially efficient if both sides of the market are perfectly competitive.*

Proof. The social planner's objective function (13) converges to the platform's profit function (3) as $\varepsilon_{D,p}$ and $\varepsilon_{H,w}$ approach infinity. □

Under perfect competition for drivers and riders (*i.e.*, all behavioral elasticities converging to infinity), the platform's pricing is first-best because the intermediary shares the same objective function and constraint as the social planner. This follows from the fact that the fractions involving behavioral elasticities in equation (13) converge to one in this situation.

This result contrasts with work that shows platform competition can be harmful (Frechette et al., 2019; Hagiu and Jullien, 2014; Tan and Zhou, 2021). The key distinction in this model is that the ratio, not the product, of agents on either side of the market governs network effects. Therefore, to the extent that platform competition increases behavioral elasticities but leaves the mapping between waiting times and utilization unaltered, greater competition brings private equilibrium outcomes closer to the socially efficient level. This feature is more appealing in contexts where market participants can multi-app and where competition comes from the threat of entry by new platforms, or from customer adoption of outside options, rather than a fracturing of agents across different platforms.[12]

Understanding socially efficient pricing in the presence of market power requires further analysis. The social planner's optimality conditions take a similar form but explicitly account for the impact of pricing changes on market participants. The social planner's first-order conditions for p and θ, respectively, are

$$1 + \tilde{\phi} \cdot \left(\varepsilon_{D,x} \cdot \varepsilon_{x,p} - \varepsilon_{D,p}\right) + \frac{\tilde{\theta}}{\frac{\varepsilon_{D,p}}{\varepsilon_{D,p}-1} - \tilde{\theta} - \tau} \cdot (1 - \varepsilon_{x,p}) = 0, \tag{14}$$

$$\varepsilon_{x,\theta} \cdot \left(\tilde{\phi} \cdot \varepsilon_{D,x} - \frac{\tilde{\theta}}{\frac{\varepsilon_{D,p}}{\varepsilon_{D,p}-1} - \tilde{\theta} - \tau}\right) = 0, \tag{15}$$

where $\phi = \frac{p \cdot (\frac{\varepsilon_{D,p}}{\varepsilon_{D,p}-1} - \theta - \tau) - c}{p \cdot (\frac{\varepsilon_{D,p}}{\varepsilon_{D,p}-1} - \theta - \tau)}$ and the notation $\tilde{\bullet}$ reflects endogenous parameters evaluated at the social optimum.

Theorem 3 (Socially efficient pricing)**.** *The socially efficient price and commission rate can be expressed as a function of elasticities that describe driver and passenger behavior*

[12] In practice, the coincidence of the platform's and the social planner's objective function under perfect competition is convenient in that it allows for a sole focus on distortions arising from market power. Further, the empirical counterfactuals below do not change the degree of competition but rather consider alternative market designs, such as commission caps and minimum wages.

as follows

$$\tilde{p} = \frac{1}{1 - \frac{1+\varepsilon_{D,x}}{\varepsilon_{D,p}}} \cdot \frac{c}{1-\tau} \cdot \frac{\varepsilon_{D,p} - 1}{\varepsilon_{D,p} + \frac{\tau}{1-\tau}}, \tag{16}$$

$$1 - \tilde{\theta} = 1 - \left(1 - \tau + \frac{1}{\varepsilon_{D,p} - 1}\right) \cdot \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}. \tag{17}$$

Proof. This follows from solving equations (16) and (17) for $\tilde{\phi}$ and $\tilde{\theta}$. □

There are two key differences with the platform's optimal solution. First, the social planner reduces passenger prices relative to the private equilibrium by multiplying p^* with the third term in equation (16). The correction does not include $\varepsilon_{D,x}$ because the platform incorporates the social planner's concern that new passengers hurt others through longer waiting times. However, the social planner corrects the monopolistic distortion and a Spence inefficiency that occurs under the private optimal pricing regime. The latter stems from the fact that the platform internalizes network effects for the marginal but not the average rider (Weyl, 2010). Second, neither of the socially efficient pricing conditions involves drivers' behavioral responses. This is a consequence of the envelope theorem since adjustments in labor supply that stem from price or commission rate changes do not have a first-order effect on driver welfare local to the optimum.

2.4 Discussion

This subsection discusses several aspects of the model outlined above.

Labor supply. The labor supply function $H(w)$ describes the number of hours drivers work on the platform. In practice, measuring labor supply to a particular platform is difficult. This is because platforms observe when workers are "online" which generally measures the hours drivers are logged on. However online status is costless to maintain because it does not obligate individuals to do anything. For example, drivers can be at home without the intention of accepting jobs but still appear online, or they may be working for a competing platform.

The concept of hours in this paper corresponds to a metric of labor supply which translates into utilized hours consistently given demand and prices. In other words, it is a structural measure of platform-specific labor supply. The platform does not need to observe this measure to price optimally. Instead, they can experiment to

reach their optimal pricing structure. Therefore, the empirical approach below circumvents the issue of directly observing platform-specific labor supply since driver supply elasticities can be inferred from platform behavior.

Pricing. The model assumes that the platform enforces a constant price and commission rate. However, platform prices and commission rates may be state-dependent. Rather than accounting for the intricacies of these high-frequency pricing strategies,[13] this model aims to provide a bird's-eye view of platform behavior that is informative of market power with minimal data requirements. This is particularly useful if platforms use a bracketing heuristic to make decisions. In other words, platforms set baseline prices and commission rates to maximize profits and then subsequently finesse their state-dependent pricing. The fact that the ratio of revenue to gross bookings in Uber's public financial filings has been so constant, as well as public comments by the platform's CEO,[14] suggests that this is approximately the case.

Costs. I model platforms as facing two variable costs τ and c, which reflect *ad valorem* and marginal costs. In addition, platforms may face fixed costs. For example, maintaining the code base and data centers that underly their services. These costs should not influence optimal pricing, which trades off marginal revenue and costs. However, if fixed costs comprise the bulk of a platform's overall costs, they may implement a reservation profit share for mediating exchanges. This can be incorporated into the model by considering marginal costs that are higher than otherwise.

Costs may also stem from attracting and maintaining riders and drivers on the platform. This would be analogous to hiring costs in models of imperfect labor market competition (Manning, 2006). But, given the digital nature of most platforms under consideration, such costs are likely subsumed into a platform's fixed cost. For instance, the same software facilitates all drivers' onboarding procedures, and riders' details are stored on the same server, where the marginal cost of storage is minimal. The fact that Uber's marketing spend and headcount are only weakly correlated with revenue growth further supports this hypothesis.

[13] See Castillo (2023) for a treatment of this phenomenon.

[14] Uber's CEO, Dara Khosrowshahi, has said, "[the platform] optimizes for an average take-rate". See the interview with The Rideshare Guy here.

3 Redesigning the Marketplace

This section considers alternative market designs to remedy platform monopsony power. Specifically, I consider two policies. First, a strategically set commission cap to raise worker welfare. Second, a minimum wage for utilized hours, as has been implemented by many state and local governments in the US (*e.g.*, most recently in Minneapolis, Minnesota).

3.1 Commission Caps and Driver Unionization

I consider a three-period game to analyze the impact of a commission cap. In *period one*, an organization sets the commission rate to maximize drivers' hourly earnings, as described in equation (1). This is equivalent to maximizing worker welfare under an isoelastic labor supply curve. Examples of such an organization would be a drivers' union with a specific mandate to control commission rates or a specially authorized government body. For convenience, I will refer to this organization as the union.

In *period two*, the platform selects the price for an hour of ridesharing services to maximize their profits. They do this with knowledge of the commission rate cap from period one and subject to the equilibrium mechanics of the marketplace summarized in equation (2), and the optimal behavior of riders and drivers as embodied in the demand and supply functions $D(\bullet)$ and $S(\bullet)$, respectively.

Finally, in *period three*, the marketplace's participants make their decisions taking the commission rate and price as given, an equilibrium is reached, and outcomes are realized. Note that workers must necessarily be better off because they can always implement the commission rate that the platform would have wanted to implement.

Backward induction solves the game between the platform and the union in the following steps. The platform's optimal choice of price given a commission rate in equation (4) implies that

$$\varepsilon_{p,\theta} = \frac{\theta}{1-\theta-\tau}. \tag{18}$$

where $\varepsilon_{p,\theta} = \frac{\partial p}{\partial \theta} \cdot \frac{\theta}{p}$. Next, I solve for the commission rate that maximizes workers' wages. The union's optimization problem is subject to two constraints. First, utilization will respond to bring the market to equilibrium, which affects wages. Second, the union internalizes the platform's optimal pricing response to changes in the commission rate with the best response function $P(\theta)$. Formally, the problem can be written

down as

$$\max_{\theta} \; p \cdot \theta \cdot x \;\text{ subject to }\; p = P(\theta) \;\text{ and }\; D(p,x) = x \cdot H(p \cdot \theta \cdot x), \tag{19}$$

which yields the first-order condition

$$1 + \varepsilon_{p,\theta} - \tilde{\varepsilon}_{x,\theta} = 0, \tag{20}$$

The term $\tilde{\varepsilon}_{x,\theta}$ differs from $\varepsilon_{x,\theta}$, which is defined in equation (6), because the union internalizes the best response of the platform in the equilibrium condition. Now, this elasticity equals

$$\tilde{\varepsilon}_{x,\theta} = \frac{\varepsilon_{D,p} \cdot \varepsilon_{p,\theta} + \varepsilon_{H,w} \cdot (1 + \varepsilon_{p,\theta})}{\varepsilon_{D,x} + 1 + \varepsilon_{H,w}}. \tag{21}$$

Theorem 4 (The union's commission rate)**.** *The commission rate that maximizes drivers' wages is*

$$1 - \theta^{**} = 1 - (1 - \tau) \cdot \left(\frac{1}{\varepsilon_{D,p}} + \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} \right). \tag{22}$$

Proof. This follows from plugging equations (21) and (18) into equation (20). □

Comparing this condition to equation (9) reveals that the wage maximizing commission rate is necessarily lower than the level set by a monopsonistic platform. Rather than marking up the commission rate according to the labor supply elasticity, the union reduces the commission rate depending on the elasticity of demand to price. The union knows the platform will increase prices if the commission rate falls. Consequently, when passengers are insensitive to prices, the union lowers the commission rate anticipating both that the platform will increase prices and that this will be largely tolerated by riders.

Theorem 5 (The optimal commission cap)**.** *The union's optimal commission rate, which is described by equation* (22)*, is the social planner's first best commission rate under perfect competition for passengers.*

Proof. Equation (22) converges to equation (17) under perfect competition for passengers. □

Theorem 5 provides a rationale for allowing drivers to unionize with a specific mandate to set the commission rate. This institutional framework restores the com-

mission rate to the social planner's desired level when the market for passengers is competitive. Further, in such a situation, the surplus for passengers is necessarily small and largely unaffected by the policy.[15] Therefore, if it were desirable to implement the policy in the first place, policymakers need not worry about the impact on riders.

Driver wages and welfare. Translating changes in driver wages to welfare effects requires two considerations. First, driver preferences. Under the assumption of an isoelastic labor supply function, drivers' surplus in this market equals

$$U(w) = \frac{w \cdot H(w)}{1 + \varepsilon_{H,w}}. \tag{23}$$

Increases in the wage rate benefit workers directly and via the number of hours worked on the platform. The change in welfare due to an exogenous change in the commission rate, which is of primary interest in the counterfactual here, is given by

$$\varepsilon_{U,\theta} = (1 + \varepsilon_{H,w}) \cdot \varepsilon_{w,\theta}, \tag{24}$$

where $\varepsilon_{U,\theta} = \frac{dU}{d\theta} \cdot \frac{\theta}{U}$. In turn, changing the rate of commission affects wages through a number of channels, as shown by the elasticity of wages to the commission rate

$$\varepsilon_{w,\theta} = 1 + \varepsilon_{P,\theta} - \tilde{\varepsilon}_{x,\theta}, \tag{25}$$

where $\varepsilon_{w,\theta} = \frac{dw}{d\theta} \cdot \frac{\theta}{w}$. Reducing the commission rate mechanically increases wages as workers keep a larger share of revenue, it raises prices due to the platform's behavioral response, which has an ambiguous effect on wages through its equilibrium consequences, and it encourages higher driver supply reducing utilization.

Second, earnings from the gig economy make up only a fraction of workers' overall income, typically around one-quarter (Anderson et al., 2021). Changes to total worker welfare can be approximated by multiplying the percentage change in worker surplus from gig work by the share of overall income earned in the gig economy (Christensen and Osman, 2023), assuming quasi-linear utility and that the elasticity of labor supply to other activities is similar.

[15] This follows from the fact that rider surplus is $\frac{p \cdot D(p,x)}{\varepsilon_{D,p}-1}$, which equals zero under perfect competition for riders.

3.2 A Minimum Wage on Utilized Hours

This subsection considers setting a minimum wage for workers' utilized hours. Such policies have been popular amongst state and local policymakers (*e.g.*, see Seattle and Minneapolis) because they do not require knowledge of workers' platform-specific labor supply. Denoting the level of this minimum wage with $\bar{w}$, such a policy ensures that $p \cdot \theta$ does not fall below $\bar{w}$.

To evaluate the impact of this policy, I study raising $\bar{w}$ marginally above the utilized wage rate $p^* \cdot \theta^*$ that prevails in the *status quo* equilibrium. The impact of this policy on wages is summarized by

$$\varepsilon_{w^*,\bar{w}} = 1 - \varepsilon_{x,\bar{w}}, \tag{26}$$

where $w^* = p^* \cdot \theta^* \cdot x$ and $\varepsilon_{x,\bar{w}} = -\frac{dx}{d\bar{w}} \cdot \frac{\bar{w}}{x}$. The elasticity of utilization to the minimum wage $\varepsilon_{x,\bar{w}}$ can be expressed in terms of behavioral elasticities after differentiating the equilibrium condition with the minimum wage substituted in

$$D\left(\frac{\bar{w}}{\theta}, x\right) = x \cdot H(\bar{w} \cdot x), \tag{27}$$

which gives

$$\varepsilon_{x,\bar{w}} = \frac{\varepsilon_{D,p} \cdot (1 - \varepsilon_{\theta,\bar{w}}) + \varepsilon_{H,w}}{1 + \varepsilon_{H,w} + \varepsilon_{D,x}}, \tag{28}$$

where $\varepsilon_{\theta,\bar{w}} = \frac{d\theta}{d\bar{w}} \cdot \frac{\bar{w}}{\theta}$. Equation (28) contains the elasticity of the commission rate to the minimum wage, which captures the platform's pricing response to the minimum wage alongside the broader equilibrium adjustments.

Characterizing the platform's reaction with behavioral elasticities requires solving their problem. The platform's optimization problem is now

$$\max_{p,\theta} \left[p \cdot (1 - \theta) - c\right] \cdot D(p, x) \text{ subject to } D(p, x) = x \cdot H(p \cdot \theta \cdot x) \text{ and } \bar{w} = p \cdot \theta. \tag{29}$$

Platform optimization implies that

$$1 - \theta^\dagger = 1 - \frac{\bar{w}}{c + \bar{w}} \cdot (1 - \tau) \cdot \frac{\varepsilon_{D,p} - \varepsilon_{D,x} \cdot \check{\varepsilon}_{x,\theta} - 1}{\varepsilon_{D,p} - \varepsilon_{D,x} \cdot \check{\varepsilon}_{x,\theta}}, \tag{30}$$

where

$$\check{\varepsilon}_{x,\theta} = \frac{dx}{d\theta} \cdot \frac{\theta}{x} = \frac{\varepsilon_{D,p}}{1 + \varepsilon_{H,w} + \varepsilon_{D,x}}, \tag{31}$$

which comes from totally differentiating equation (27) with respect to x and θ. The dagger notation denotes the platform's endogenous choices in this new environment. Solving for prices readily follows from substituting equation (30) into the minimum wage constraint $p^{\dagger} = \frac{\bar{w}}{\theta^{\dagger}}$.

Theorem 6 (Worker welfare improving minimum wage on utilized hours.)**.** *The minimum wage on utilized hours leaves workers better off if and only if*

$$c + \bar{w} < \frac{1}{1 - \frac{1+\varepsilon_{D,x}}{\varepsilon_{D,p}}}. \tag{32}$$

Proof. This follows from setting equation (26) greater than zero, recognizing that $\varepsilon_{\theta,\bar{w}} = \frac{1}{c+\bar{w}}$ from equation (30), and substituting in. □

Unlike a commission cap, a minimum wage on utilized hours is not necessarily beneficial for workers. However, it is more likely to be so the when platform's market power over passengers is strong. The righthand side of inequality (32) equals the platform's markup on rider prices above marginal cost. Intuitively, platforms would rather increase the price than give a larger share of revenue to drivers. Therefore, when the platform's monopoly power is high, price increases do not reduce utilization too much, and drivers are left better off at the expense of riders.

4 An Application to Uber

In this section, I use the model from section 2 to evaluate the extent of monopsony power enjoyed by the US's largest ridesharing platform, Uber. The model's parsimonious structure facilitates this analysis using only publicly available data and causal estimates from the academic literature. The results suggest that Uber enjoys substantial monopsony power over drivers but faces a competitive market for riders. Setting the commission rate to its first-best level, or equivalently, empowering a commission rate setting union, increases wages by 14 percent. Despite this, a minimum wage on utilized hours likely harms workers.

4.1 Institutional Details

Uber was founded in 2009 and has grown to operate in 72 countries globally. It is the largest ridesharing platform in the US, with an estimated market share of around 75 percent. Currently, Uber has 1.5 million earners on its platform in the US. In most areas, workers are free to join and leave the platform, and once they are on the platform, drivers pick where and when to work.[16] Drivers can also work simultaneously for Uber's competitors, like Lyft.

Given the available data, the focus of the analysis in this paper is Uber's US ridesharing marketplace around 2017. During this time, passenger fares were determined by time, distance, and Uber's surge algorithm. Two components determined fares: the price of the ride and a booking fee. Drivers on the platform received the price component of the fare after the Uber fee, which was a fixed rate, was deducted. All the booking fees went to Uber to cover the costs of mediating the ride. Tipping was only introduced in mid-2017 and was very rare (Chandar et al., 2019).

4.2 Data

Estimation requires three empirical moments, and information on Uber's costs, to identify the model's three structural parameters: $\varepsilon_{D,p}$, $\varepsilon_{D,x}$, and $\varepsilon_{H,w}$. Uber's commission rate (*i.e.*, $\widehat{1-\theta}$) provides the first empirical moment. This number is the subject of significant discussion, which is often confused by the coexistence of the booking fee for passengers and the Uber fee for drivers. However, the model provides a clear theoretical definition of the commission rate: the share of the total price paid by riders—inclusive of the booking fee—that drivers do not receive. Therefore, information on the average passenger fare, booking fee, and Uber fee is necessary to construct an estimate of the commission rate.

I take these numbers from academic publications that have access to proprietary microdata, and cross-check their implications with public sources like online Uber driver fora. Recent papers report an Uber fee ranging from 20 to 28 percent (Caldwell and Oehlsen, 2021; Castillo, 2023; Cook et al., 2021). In the estimation, I use an Uber fee of 25 percent for the central scenario, which seems to be Uber's active choice for the commission rate in 2017.[17] In an earlier working paper from 2019, Castillo (2023) reports a booking fee of $2.30 for Houston, Texas. This is on the higher side of reports

[16] A notable exception is New York, where the city has regulated limits on the onboarding of drivers.

[17] Some drivers had a lower Uber fee in that year because they were grandfathered in from previous regimes.

of the booking fee from drivers during that period of time,[18] so I opt for a lower booking fee of $1.30 to calculate the overall commission rate. Finally, Cook et al. (2021) reports drivers' earnings per trip before the Uber fee, which, when combined with the booking fee, implies an average price per trip of $11.40.

Overall, these numbers constitute a commission rate of 34 percent, which I use as the central scenario in the analysis below.[19] I also consider commission rates of 29 percent and 39 percent. As well as reflecting some uncertainty about the true value of the commission rate, these numbers are also indicative of where Uber's commission rate used to be before 2017, when the platform was more generous to drivers, and where the commission rate is suggested to be at present after recent pricing changes.

The second empirical moment is the price Uber charges for an hour of ridesharing services. Combining the average price per trip of $11.40 with the average length of a trip produces this number. Fortunately, Cook et al. (2021) reports the average trip speed and distance, which jointly suggest a typical length of just over fifteen minutes. In turn, this suggests a price of $43.59 for one hour's worth of ridesharing services.

Theoretically, behavioral elasticities and costs comprise the price Uber charges. Consequently, information on Uber's costs is also required for the estimation. The main marginal costs to mediating exchanges are transaction fees for payment processing, sales tax payable to local government, and insurance coverage for drivers against "life-changing events". Again for Texas, Houston, Castillo (2023) reports the first two components comprise three percent of the fare. Insurance costs are paid by the mile at an approximate premium of $0.30. Combined with the average trip distance, inclusive of distance to pick up, this suggests that insurance costs make up 15 percent of the passenger fare. In total, costs comprise 18 percent of the typical fare. To examine the sensitivity of estimates to uncertainty in this calculation, I also consider total costs equivalent to 13 percent and 23 percent of the fare. I assume these stem from changes in insurance costs, which have been volatile over time.

The third and final empirical moment is the equilibrium response of utilization to a change in price.[20] Hall et al. (2023) report static and dynamic estimates of this statistic, which exploit pricing experiments by the platform. Given that base pricing is driven by long-term considerations, I use the dynamic estimate, which is six months out from the price change, and its standard error from Figure 5 in the paper. I infer

[18] See discussion here.

[19] The choice of commission rate is also supported by Uber's breakdown of gross bookings in this blog post.

[20] The equilibrium response of utilization to a change in the commission rate would provide an overidentifying restriction but, unfortunately, I am not aware of any estimates of this statistic.

a central estimate of 1.40 with a standard error of 0.38 (= 0.75/1.96). This estimate comes from price experiments in several large US cities between 2014 and 2017.

The measure of utilization in this empirical moment uses online hours in the denominator, which differs from the relevant concept of platform-specific labor supply. To correct for this, I leverage the structure of the model to adjust the measure of utilization during the estimation. This makes use of a further moment that is reported in Hall et al. (2023), namely, the elasticity of online hours to earnings $\hat{\varepsilon}_{H,w}$ (= 6.39) and the following Taylor series approximation

$$\varepsilon_{x,p} \approx \hat{\varepsilon}_{x,p} + \frac{\partial \varepsilon_{x,p}}{\partial \varepsilon_{H,w}} \cdot (\varepsilon_{H,w} - \hat{\varepsilon}_{H,w}) = \hat{\hat{\varepsilon}}_{x,p}, \tag{33}$$
$$\text{where} \quad \frac{\partial \varepsilon_{x,p}}{\partial \varepsilon_{H,w}} = \frac{1 - (\varepsilon_{D,p} - \varepsilon_{D,x})}{(\varepsilon_{D,x} + 1 + \varepsilon_{H,w})^2},$$

where $\hat{\varepsilon}_{x,p}$ is the elasticity of utilization with respect to price measured using online hours. So $\hat{\hat{\varepsilon}}_{x,p}$ is used as the third empirical moment in the estimation. In practice, this does not impact estimates noticeably.

Combining the numbers above with further data on the average number of trips per week, hours per week, and driving speed from Cook et al. (2021) implies other interesting numbers. In particular, they suggest an average wage of $14.72, a utilization rate of 51 percent,[21] and a utilized wage rate of $28.96. This is on the high side of Uber's reported earnings per utilized hour, which suggests that the statistics above do not offer a particularly negative picture of drivers' earnings.[22]

4.3 Estimation

I use a generalized method of moments estimator to estimate the model's structural parameters.[23] Precisely, I select $\boldsymbol{\varepsilon} = (\varepsilon_{D,p}, \quad \varepsilon_{D,x}, \quad \varepsilon_{H,w})$ to minimize the distance between $\hat{X} = (\widehat{1-\theta}, \quad \hat{p}, \quad \hat{\hat{\varepsilon}}_{x,p})$ and the model's predictions from equations (6), (8),

[21] This utilization rate only includes time with passengers and corresponds to x in the model.

[22] See this blog post again.

[23] This exercise is close to calibration, however, the standard error from Hall et al. (2023) for the elasticity of utilization to prices allows the estimation procedure to quantify the statistical uncertainty stemming from this moment. Although this standard error is large, the estimates of the behavioral elasticities from the model are precise.

	Data moment	Model prediction
Commission rate	0.34	0.34
Price	43.59	43.53
Utilization elasticity	1.4 [0.65, 2.15]	1.18

Table 1: Model Fit

Notes: This table shows the targeted moments in the first column, their empirical estimates in the second column, and the model's predictions of these moments in the third column. The numbers in the parentheses are the 95 percent confidence interval for the empirical estimate of the utilization elasticity.

and (9) using the norm $m(\hat{X},\varepsilon)^T \cdot W \cdot m(\hat{X},\varepsilon)$, where

$$m(\hat{X},\varepsilon) = \begin{pmatrix} \widehat{(1-\theta)} - 1 - (1-\tau)\cdot\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}\cdot\frac{\varepsilon_{H,w}}{1+\varepsilon_{H,w}} \\ \hat{p} - \frac{1}{1-(\frac{1}{\varepsilon_{D,p}}+\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}})}\cdot\frac{\hat{c}}{1-\hat{\tau}} \\ \hat{\hat{\varepsilon}}_{x,p} - \frac{\varepsilon_{D,p}+\varepsilon_{H,w}}{\varepsilon_{D,x}+1+\varepsilon_{H,w}} \end{pmatrix}, \tag{34}$$

and W is the weighting matrix. Although the model is just identified, weighting is helpful because of the finite sample.

I produce standard errors for the estimates by sampling 500 values of $\hat{\varepsilon}_{x,p}$ from a normal distribution with a mean of 1.40 and a standard deviation equal to 0.38. Therefore, these standard errors reflect only statistical uncertainty from the empirical estimate of the elasticity of utilization to price. The sensitivity of results to the commission rate and markup is checked by re-estimating the parameters under different assumptions about these moments.

Table 1 compares the model's predictions with the baseline empirical moments. The model fits the three data moments extremely well. Although this is unsurprising since the model is exactly identified, it is not completely trivial because of the finite sample and sign restrictions on the elasticities. Further, other empirical models have not been able to reconcile Uber's behavior purely with profit maximization (Castillo, 2023; Rosaia, 2020).

Differences between short- and long-run behavioral elasticities can explain this issue. Short-run elasticities that exploit variation in surge pricing, or experiments

that last less than a few weeks, are generally small. Therefore, when they are used to compute passenger and driver behavioral responses, these agents appear inelastic, which suggests that Uber has a lot of market power and should charge higher prices. Other long-term pricing experiments on the Uber platform have found much larger elasticities (Christensen and Osman, 2023).

4.4 Seattle's *Fare Share* Ordinance

Another way to evaluate the model is to test its out-of-sample performance. In this subsection, I compare the fallout of Seattle's *Fare Share* ordinance, which came into force at the start of 2021, with the model's predictions.[24] This regulation effectively placed a minimum wage on workers' utilized hours by imposing minimum levels of payments to drivers based on a trip's distance and duration. At the time, drivers were required to receive at least $1.33 per mile and $0.57 per minute, or a minimum of $5.00 per trip.[25] In response, Uber raised prices by 40 percent.[26]

Interpreting this through the model, it is possible to map Uber's price response to how the policy affected drivers' utilized earnings. Equation (30), in combination with the minimum wage constraint, implicitly describes the platform's optimal price. The elasticity of prices to the minimum wage on utilized hours equals 0.90, when evaluated at the calibrated level of costs and existing utilized wage rate. Therefore, Uber's price response indicates the policy raised utilized wages by 45 percent. Uber do not report this number for Seattle, but the platform estimates that its labor costs will rise by up to 40 percent in the face of similar proposals in Minnesota,[27] which are less tough than those for Seattle at the time.

If the ordinance raised utilized wages by 45 and prices increased by 40 percent, then the platform would have to raise commission rates by five percent or, equivalently, three percentage points. A small increase in commission rates is consistent with the model, which predicts that Uber would respond to the policy primarily through price adjustments rather than changes in commission. Taking the price and commission rate changes together, utilization would fall by 56 percent, causing overall wages to fall by 11 percent. Uber reports that wages per online hour fell by ten percent,[28]

[24] I do not use this event to provide over-identifying restrictions for the estimation because it occurred four years after the other data moments.

[25] This has since been superseded by state-level legislation that requires at least $1.55 per mile and $0.66 per minute, or $5.81 per trip

[26] See this Uber blog post.

[27] See this Uber blog post

[28] See this Uber blog post again.

		Commission rate		
		39%	34%	29%
Costs	13%	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.90$ (<0.01) $\varepsilon_{H,w} = 2.36$ (<0.01)	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.90$ (<0.01) $\varepsilon_{H,w} = 3.23$ (0.01)	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.90$ (<0.01) $\varepsilon_{H,w} = 4.39$ (0.01)
	18%	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.84$ (<0.01) $\varepsilon_{H,w} = 2.92$ (0.01)	$\mathbf{\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.84}$ (<**0.01**) $\mathbf{\varepsilon_{H,w} = 4.27}$ (**0.02**)	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.84$ (<0.01) $\varepsilon_{H,w} = 6.37$ (0.03)
	23%	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.79$ (<0.01) $\varepsilon_{H,w} = 3.84$ (0.02)	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.79$ (<0.01) $\varepsilon_{H,w} = 6.31$ (0.03)	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.79$ (<0.01) $\varepsilon_{H,w} = 11.63$ (0.08)

Table 2: Parameter Estimates

Notes: This table shows a matrix of parameter estimates for nine different combinations of Uber's commission rate and costs. Left to right shows increasingly lower commission rates. Up to down shows increasingly higher costs. Parentheses show corresponding standard errors. The estimates in the central cell in bold are the central scenario.

again, matching the model's prediction closely.

The takeaways from this subsection are twofold. First, the model accurately predicts the response of prices and equilibrium outcomes to policy interventions out-of-sample. Second, Seattle's minimum wages on utilized hours does not seem to have benefited workers. I evaluate the efficacy of these policies in the US more generally in section 5.2.

4.5 Results

Table 2 shows parameter estimates for nine different combinations of Uber's commission rate and costs. The central scenario is highlighted in bold in the middle of the matrix. All of these variations find that Uber faces a very competitive market for riders (*i.e.*, high values of $\varepsilon_{D,p}$ and $\varepsilon_{D,x}$) so, for ease of interpretation, I report the ratio of the elasticity of demand to utilization and price $\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}$. This approximately equals the driver's keep rate (*i.e.*, one minus the commission rate) under a perfectly competitive driver market.[29] The fact that all these ratios all but equal one minus the cost share under consideration confirms the highly competitive rider market; the commission rate would only cover costs were it not for the platform's ability to markup thanks to

[29] When multiplied by $(1-\tau)$, which is close to one, this mapping is exact.

monopsony power.

In contrast, the results suggest that Uber exerts significant market power over drivers. The central estimate, which is highlighted in bold at the center of table 2, implies that the platform faces a driver supply elasticity of 4.27. This number is remarkably similar to estimates of monopsony power in other US labor markets despite the very different modeling and estimation approach (Lamadon et al., 2022). The estimate of monopsony power decreases if the platform is considered to charge a higher commission rate and rises if Uber is believed to face higher costs. All the estimates imply a considerable degree of monopsony power unless one maintains that Uber has *both* an unlikely high level of costs and low commission rate.

In a standard model of wage-posting by a monopsonistic employer, the driver supply elasticities map directly to a wage markdown of $1/(1+\varepsilon_{H,w})$. For the central estimate, this implies that workers would be denied one-fifth of their marginal product. In the two-sided market described in section 2, this is not the case because equilibrium adjustments in utilization determine wages. In section 5, I explore the impact of Uber's monopsony power on wages and welfare by considering feasible counterfactuals that account for equilibrium effects.

5 Counterfactuals

This section studies two counterfactuals to quantify how monopsony power affects wages and, in turn, worker welfare in a two-sided ridesharing market. Specifically, I study a commission cap set to maximize driver welfare and a minimum wage for workers' utilized hours.

5.1 Commission Caps and Driver Unionization

Motivated by Theorem 5, I consider setting the commission rate equal to its socially efficient level in equation (17) and allowing the platform to respond with rider prices. This is equivalent to a union selecting a commission rate that maximizes wages when there is competition for riders, which table 2 evinces in the case of Uber. A practical advantage of this formulation is that it leaves the denominator in the welfare expression (23) constant, and the platform's pricing adjustment is straightforward.

Using equations (24) and (25), table 3 presents estimates of the impact of monopsony power on wages w and the worker surplus solely derived from Uber's market-

Costs	Commission rate: 39%	34%	29%
13%	$\%\Delta w = 34$	$\%\Delta w = 23$	$\%\Delta w = 16$
	$\%\Delta U = 114$	$\%\Delta U = 98$	$\%\Delta U = 84$
18%	$\%\Delta w = 23$	$\%\mathbf{\Delta w = 14}$	$\%\Delta w = 8$
	$\%\Delta U = 92$	$\%\mathbf{\Delta U = 74}$	$\%\Delta U = 58$
23%	$\%\Delta w = 14$	$\%\Delta w = 7$	$\%\Delta w = 2$
	$\%\Delta U = 70$	$\%\Delta U = 50$	$\%\Delta U = 31$

Table 3: Welfare Effects of a Commission Cap

Notes: This table shows a matrix of estimates for changes in Uber's wage $\%\Delta w$ and worker surplus $\%\Delta U$ estimates for nine different combinations of Uber's commission rate and costs. Left to right shows increasingly lower commission rates. Up to down shows increasingly higher costs. The estimates in the central cell in bold are the central scenario.

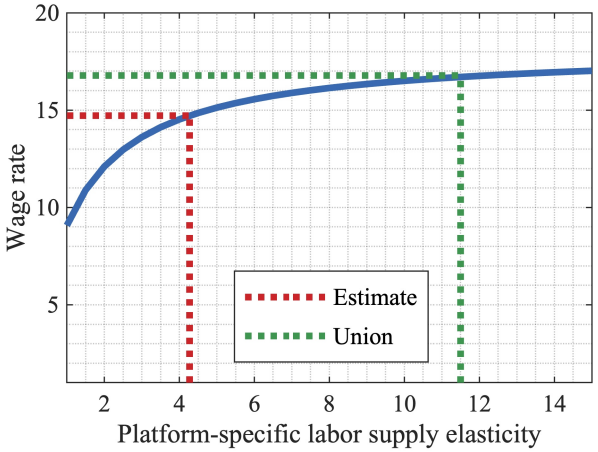
place U in terms of percentage changes. The central estimate in bold implies that drivers' wages would rise by 14 percent in equilibrium, or $2.00 per hour. It is possible to decompose this change in wages using equation (25) as follows

$$\%\Delta w = \left[1 + \underbrace{(1-\varepsilon_{x,p})\cdot\varepsilon_{P,\theta}}_{\text{(1-1.18)}\times\text{2.17=-0.40}} - \underbrace{\varepsilon_{x,\theta}}_{\sim 0}\right] \cdot \underbrace{\%\Delta\theta}_{0.23} \approx 14. \tag{35}$$

Pricing responses by the platform and equilibrium adjustments in utilization mediate the effect of changes in the commission rate on wages. The elasticity of the platform's price to the driver's keep rate is 2.17, as computed from equation (18). This has a further positive effect on drivers' wages *ceteris paribus*. However, the increase in prices also triggers an equilibrium adjustment in utilization. This equilibrium response outweighs the positive effect on wages from the platform raising prices because $1-\varepsilon_{x,p}$ is negative. Reducing commission rates also decreases utilization further, although the impact of this is approximately zero because the rider market is so much more competitive than the driver market.

The range of wage effects varies predictably with the extent of the platform's monopsony power. The highest estimate implies that wages are almost one-third below their counterfactual equivalent. At the lower end, wages are only minimally affected by a

Figure 1: Wages as a Function of the Platform-Specific Labor Supply Elasticity



Notes: This figure plots drivers' equilibrium wage rate in solid blue as a function of Uber's platform-specific labor supply elasticity, keeping demand-side elasticities constant. The dashed red line denotes the *status quo* equilibrium level of wages, and the dashed green line shows the level of wages attained in the commission cap counterfactual.

small amount of monopsony power but this scenario requires a low level of commission, which Uber no longer offers, and a high level of costs. Taken together, the evidence suggests that the platform materially depresses wages relative to the counterfactual. However, these estimates are lower than other papers that combine short-term variation in driver earnings to estimate supply elasticities with traditional wage-posting models (Caldwell and Oehlsen, 2021), and they incorporate the attenuating effect of fare and utilization adjustments.

One way to benchmark these wage changes is to compare their magnitude with equivalent changes in Uber's platform-specific labor supply elasticity. Figure 1 illustrates this exercise. Attaining the wage gains that occur under the commission cap scenario requires almost a tripling of Uber's platform-specific labor supply elasticity, which would entail a dramatic change in the competitive landscape of the US ridesharing market.

Table 3 also reports the overall effect on the ridesharing worker surplus of these wage changes, which rely on the assumption of an isoelastic driver supply function. This counterfactual leads to large welfare gains for workers. A 14 percent increase in wages raises the worker surplus from Uber's marketplace by close to three-quarters. Translating these welfare effects in the ridesharing market to overall welfare requires accounting for gig workers' total level of income. Multiplying by the share of income

that workers derive from ridesharing completes this mapping under the assumptions that utility is quasi-linear and the labor supply elasticity to other markets is the same as to ridesharing, which seems reasonable given similar estimates for other labor markets (Lamadon et al., 2022). Somewhere on the order of one-quarter of gig economy participants' income is earned through ridesharing, which suggests that the counterfactual could raise overall welfare for Uber's US drivers by almost one-fifth.

5.2 A Minimum Wage on Utilized Hours

In terms of a minimum wage on utilized hours, estimates of the model's parameters and the prevailing average wage level suggest that this policy harms workers. Evaluating the left-hand side of inequality (32) at the *status quo* utilized wage and costs level equals 37, which exceeds the right-hand side of 6. This indicates there is no room to raise utilized wages in a way that increases equilibrium wages because they would trigger a fall in utilization, which more than offsets the positive direct effect on equilibrium wages. This is exemplified by the discussion of Seattle's *Fare Share* ordinance in section 4.4.

In summary, the minimum wage is ineffective despite Uber's significant monopsony power. This type of minimum wage allows the platform to select its optimal price and commission rate mix while satisfying the minimum wage. The additional flexibility relative to a commission cap leaves the platform able to exploit its monopsony power, which can manifest in low utilization, as well as low utilized wages. Ultimately, the policy fails to target the welfare-relevant quantity: equilibrium wages.

6 Conclusion

This paper develops a tractable model of a two-sided ridesharing marketplace. The framework reveals how platforms exploit monopsony power over drivers by marking up commission rates according to the driver supply elasticity that they face. Consequently, descriptions of platform-specific labor supply are an appropriate way to measure monopsony power in these settings. However, the nature of two-sided markets complicates the final effect on workers' wages and welfare.

Redesigning these marketplaces can raise worker welfare and efficiency in ways reminiscent of one-sided labor markets. In the presence of monopsony power, allowing a union of workers to set commissions—while preserving the platform's power

to set prices—delivers the first-best commission rate if the other side of the market is sufficiently competitive. This can benefit workers even when minimum wages on utilized hours cannot.

Taking the theory to the data with publicly available information on Uber's marketplace suggests that the US's largest ridesharing platform enjoys substantial monopsony power over workers. Because the platform faces a competitive rider market, a drivers' union would set the first-best commission rate. If this were the case, commission rates would fall by 15 percentage points, wages would increase by 14 percent, and worker welfare would increase dramatically. Conversely, there is no room for minimum wages on utilized hours to benefit workers, which indicates this has been a misguided policy in many US cities and states.

References

Anderson, M., McClain, C., Faverio, M., and Gelles-Watnick, R. (2021). The state of gig work in 2021. *Pew Research Center*, 8.

Caldwell, S. and Oehlsen, E. (2021). Gender differences in labor supply: Experimental evidence from the gig economy. *Unpublished*.

Castillo, J. C. (2023). Who benefits from surge pricing? *Available at SSRN 3245533*.

Chandar, B., Gneezy, U., List, J. A., and Muir, I. (2019). The drivers of social preferences: Evidence from a nationwide tipping field experiment. Technical report, National Bureau of Economic Research.

Chen, M. K., Rossi, P. E., Chevalier, J. A., and Oehlsen, E. (2019). The value of flexible work: Evidence from uber drivers. *Journal of political economy*, 127(6):2735–2794.

Christensen, P. and Osman, A. (2023). The demand for mobility: Evidence from an experiment with uber riders. Technical report, National Bureau of Economic Research.

Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of uber. Technical report, National Bureau of Economic Research.

Cook, C., Diamond, R., Hall, J. V., List, J. A., and Oyer, P. (2021). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies*, 88(5):2210–2238.

Cullen, Z. and Farronato, C. (2021). Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms. *Management Science*, 67(7):3985–4003.

Datta, N., Rong, C., Singh, S., Stinshoff, C., Iacob, N., Nigatu, N. S., Nxumalo, M., and Klimaviciute, L. (2023). Working without borders: The promise and peril of online gig work.

Dube, A., Jacobs, J., Naidu, S., and Suri, S. (2020). Monopsony in online labor markets. *American Economic Review: Insights*, 2(1):33–46.

Fisher, J. (2022). Worker welfare in the gig economy. *Working Paper*.

Frechette, G. R., Lizzeri, A., and Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, 109(8):2954–2992.

Garin, A., Jackson, E., Koustas, D. K., and Miller, A. (2023). The evolution of platform gig work, 2012-2021. Technical report, National Bureau of Economic Research.

Hagiu, A. and Jullien, B. (2014). Search diversion and platform competition. *International Journal of Industrial Organization*, 33:48–60.

Hall, J. V., Horton, J. J., and Knoepfle, D. T. (2023). Ride-sharing markets re-equilibrate. Technical report, National Bureau of Economic Research.

Harris, S. D. and Krueger, A. B. (2015). *A Proposal for Modernizing Labor Laws for Twenty-First-Century Work: The" Independent Worker"*. Brookings Washington, DC.

Hyman, L., Groshen, E. L., Litwin, A. S., Wells, M. T., Thompson, K. P., and Chernyshov, K. (2020). Platform driving in seattle.

Jullien, B., Pavan, A., and Rysman, M. (2021). Two-sided markets, pricing, and network effects. *Handbook of Industrial Organization*.

Kingman, J. F. (1961). The single server queue in heavy traffic. *Mathematical Proceedings of the Cambridge Philosophical Society*.

Kroft, K., Luo, Y., Mogstad, M., and Setzler, B. (2020). Imperfect competition and rents in labor and product markets: The case of the construction industry. Technical report, National Bureau of Economic Research.

Lachowska, M., Mas, A., Saggio, R., and Woodbury, S. A. (2023). Work hours mismatch. Technical report, National Bureau of Economic Research.

Lamadon, T., Mogstad, M., and Setzler, B. (2022). Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, 112(1):169–212.

Lerner, A. P. (1934). Economic theory and socialist economy. *The Review of Economic Studies*, 2(1):51–61.

Manning, A. (2006). A generalised model of monopsony. *The Economic Journal*, 116(508):84–100.

Manning, A. (2011). Imperfect competition in the labor market. In *Handbook of labor economics*, volume 4, pages 973–1041. Elsevier.

Prassl, J. (2018). *Humans as a service: The promise and perils of work in the gig economy*. Oxford University Press.

Ravenelle, A. J. (2019). *Hustle and gig: Struggling and surviving in the sharing economy*. Univ of California Press.

Robinson, J. (1969). *The economics of imperfect competition*. Macmillan Company.

Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029.

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.

Rosaia, N. (2020). Competing platforms and transport equilibrium: Evidence from new york city. *Working Paper*.

Rysman, M. (2009). The economics of two-sided markets. *Journal of economic perspectives*, 23(3):125–143.

Sullivan, M. (2022). Price controls in a multi-sided market. Technical report, Working paper.

Tan, G. and Zhou, J. (2021). The effects of competition and entry in multi-sided markets. *The Review of Economic Studies*, 88(2):1002–1030.

Van Reenen, J. (2024). Labor market power, product market power and the wage structure: a note. *Working Paper*.

Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.