Monopsony Power in the Gig Economy

Jack Fisher *

March 12, 2025

Abstract

Many workers provide services for customers via digital platforms that may exert monopsony power. Typical expositions of labor market power are inapplicable in this context because platforms post prices to both sides of a two-sided market instead of setting wages. Further, platform-specific labor supply is hard to measure when workers multi-app. This paper develops a model of a gig labor market that resolves these issues. Platforms exploit monopsony power to markup their commission rate, reduce equilibrium wages, and do not lower prices for customers. Estimating the model with public data on Uber implies that the platform uses labor market power to depress drivers' earnings by 15 percent. Commission caps are an effective policy to raise worker welfare, while minimum wages on utilized hours, which are common, likely harm workers.

^{*}University of Virginia: jack.w.fisher.910gmail.com. I am thankful to participants at the NBER's Digital Economics Meetings and the TSE Digital Economics Conference, as well as Johsua Angrist, Richard Blundell, Matteo Bobba, Sydnee Caldwell, Florian Ederer, Chiara Farronato, John Horton, Myrto Kalouptsidi, Matthew Leisten, Michael Luca, 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 hundreds of millions of workers around the world provide short-term and independent labor services via digital intermediaries (Datta et al., 2023; Dube et al., 2020).¹ Concerns around monopsony power in the gig economy are particularly pronounced given workers' self-employed status, which offers few protections, and fears over poor outside options owing to, for example, underemployment (Lachowska et al., 2023). Yet, policymakers lack a framework describing how monopsony power manifests in the gig economy, where platforms post prices and commission rates rather than wages, and what remedies may be effective.

This paper develops a tractable model to study a typical gig labor market: ridesharing. Alongside food delivery, this industry represents 90 percent of the over five million platform workers in the US (Garin et al., 2023). The model provides clear insights into platform pricing, the merits of different policy interventions, and their interplay with market power from both a platform and a social planner's perspective. Estimating the model requires only a small number of statistics and circumvents the need to measure drivers' platform-specific hours, which is difficult when workers multi-app and are not obliged to accept rides (Hyman et al., 2020).²

I demonstrate the framework's utility by testing its out-of-sample predictions and evaluating the extent of market power enjoyed by the US's largest ridesharing platform, Uber.³ Public data, including causal estimates from the platform's pricing experiments, indicate substantial monopsony power over drivers and a competitive

¹An 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.

²This measurement issue has prevented the implementation of a minimum wage for gig workers (Harris and Krueger, 2015).

³Bloomberg estimates the platform accounts for 75 percent of the US ridesharing industry.

market for passengers. The analysis suggests Uber 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 because this would increase congestion further. Consequently, monopsony power hurts both drivers and passengers—a key legal test for anti-trust in two-sided markets.⁴

Quantitatively, the commission rate is 15 percentage points higher than the social planner's optimum. If this were restored to its first-best level via a commission cap, wages would increase by 14 percent after accounting for the platform's pricing response. This captures almost all of the increase in wages that would occur under perfect competition. In contrast, minimum wages on utilized hours likely harm drivers. These policies ensure minimum payments to drivers for the time they spend carrying passengers. To meet this constraint, platforms increase prices rather than decrease commission rates, which dampens demand and, in turn, lowers equilibrium wages.

This has practical policy implications. State and local governments already enact minimum wages on utilized hours, which entail regulating a combination of commission rates and prices, so commission caps are feasible as well as effective. The model highlights two additional benefits of commission caps as a response to monopsony power in gig labor markets. First, given the power to set commission rates, policymakers can induce variation sufficient to identify the optimal policy. Second, if the passenger market is competitive, then drivers would collectively set the first-best commission rate given the option. Therefore, in this scenario, policymakers could ask workers to inform their decision-making without fears of strategic reporting.

Concretely, the model considers a ridesharing platform 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 various micro-foundations. Hourly wages, which also depend on utilization since drivers are only paid when they carry passengers, determine the supply of drivers. The market reaches equilibrium through adjustments in worker utilization; drivers enter and exit as their wage rate moves with utilization, and riders change their demand as waiting times fluctuate accordingly (Hall et al., 2023). This implies a fixed point equilibrium condition that constrains the platform's problem.⁵

⁴See the US's Supreme Court opinion in Ohio v. American Express Co. (06/25/2018).

⁵The 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 markets that exhibit congestion, which is common in the gig economy (*e.g.*, food delivery). 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 alternatives outside of the gig economy.

The model delivers clear intuitions about decision-making by digital 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 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 platformspecific 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 passenger prices and driver utilization, are necessary to infer the equilibrium impact on wages.

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. 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. Therefore, both sides of the market are left worse off by the presence of labor market power.

The small number of parameters in the model makes it easy to estimate with little information, which is especially valuable in a context where data is proprietary and collaborations on contentious topics are not feasible. Further, by virtue of inferring the model's parameters from platform choices, any empirical analysis is readily reconcilable with profit maximization and suitable for counterfactual evaluation. This has not been possible in other studies of multi-sided transport markets that provide a more detailed description of participants' interactions but yield behavioral responses inconsistent with standard firm objectives (Castillo, 2023; Rosaia, 2020).

To illustrate the model's utility, I evaluate the extent of Uber's market power over the US ridesharing industry—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 (Chen et al., 2019; Fisher, 2022), concerns remain about the welfare of individuals subject to these work arrangements (Prassl, 2018; Ravenelle, 2019). Therefore, quantifying monopsony power over workers in ridesharing markets and potential remedies is a first-order policy question.

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 pricing varies from ride to ride, in the words of Uber's CEO, Dara Khosrowshahi, "[the platform] optimizes for an average take-rate".⁶ 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 behavioral elasticities that Uber faces, I use public information on the platform's choice of commission rate and passenger prices around 2017, as well as data on costs and results from a randomized pricing experiment. These numbers imply a commission rate of over one-third, mediation costs equivalent to 18 percent of the fare (Castillo, 2023; Cook et al., 2021), and a large negative response of utilization to price increases. The latter statistic comes from Hall et al. (2023), which estimates the response of utilization to randomized price changes over a six-month time horizon. The model matches the data closely and accurately predicts out-of-sample the impact of a minimum wage on utilized hours in Seattle, validating the structural assumptions that support counterfactual analysis.

Bringing the model to the data suggests that Uber exerts substantial monopsony power over drivers and faces strong competition for riders. The central scenario implies the platform faces a labor supply elasticity of 4.27.⁷ Viewed through the model, this suggests the commission rate is 15 percentage points above the competitive benchmark, which corresponds to a wage markdown of one-fifth in standard wage-posting models. However, in gig labor markets, prices and utilization are endogenous and also determine wages. Therefore, it is necessary to consider a precise counterfactual to understand the impact of monopsony power on wages and worker welfare.

I consider introducing a commission cap set at the first-best level as a potential remedy to monopsony power, noting its precedent and feasibility. In this scenario, a commission rate fall of 15 percentage points triggers the platform to raise prices by almost half. In turn, utilization falls by two-thirds so that, overall, wages rise by

⁶See Dara's interview with The Rideshare Guy here. The "take rate" is another phrase for commission. ⁷This estimate is close to, but smaller than, other firm-level estimates found across US labor markets in Lamadon et al. (2022). The similarity is noteworthy given the very different estimation approach.

14 percent. The efficacy of this policy stands in contrast to the impact of minimum wages for utilized hours, which are prevalent in the US. The model shows that under plausible parameter values, these policies lead platforms to increase prices much more than commission rates. This causes a fall in utilization and, overall, driver wages decrease—opposite to the intention of the policy.

Related literature. This paper contributes to three literatures. First, there is a large and growing body of work evincing the existence of employer monopsony power in different labor markets (see Azar and Marinescu (2024); Caldwell et al. (2024) and Manning (2021) for a review of empirical work, and Kline (2025) for an overview of theoretical treatments). This paper contributes a tractable framework to expand this analysis to multi-sided labor markets and, to my knowledge, provides the first theoretically grounded estimates of monopsony power in the gig economy. Notably, the two-sided nature of this study connects it with recent work evaluating the interaction of product and labor market power (Kroft et al., 2020; Van Reenen, 2024).

Second, this paper adds to the extensive literature on minimum wages in traditional labor markets (Dube (2019); Berger et al. (2025); Horton (2025); Neumark and Shirley (2022); Vergara (2023), to name some recent work), which closely relates to a broader body of work on pricing regulations (*e.g.*, in the rental market (Diamond et al., 2019; Glaeser and Luttmer, 2003) and in credit card markets (Rysman, 2007)). This paper derives conditions under which commission caps and minimum wages on utilized hours can improve worker welfare in the gig economy. Again, to my knowledge, these conditions provide the first formalization of intuitions motivating active policy interventions in several US cities and states and empirical analyses of commission caps (*e.g.*, in Sullivan (2022) and Li and Wang (2024)).

Third, this paper builds on empirical and theoretical research in two-sided markets more broadly (Jullien et al., 2021; Rochet and Tirole, 2003; Rysman, 2009). Empirically, it estimates a model of a two-sided market that is reconcilable with platform profit maximization and, thus, amenable to considering counterfactual pricing responses. This contrasts with richer models of multi-sided markets that keep prices fixed in counterfactuals or introduce non-structural parameters into platforms' objective functions (Lee, 2013; Yu, 2024). Theoretically, this paper contributes a tractable model of a two-sided marketplace with congestion and multi-homing (Belleflamme and Toulemonde, 2009; Karle et al., 2020). The model's parsimony yields novel insights into asymmetric seesaw effects and may prove useful for pedagogical purposes. 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.

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 hour of ridesharing services p and driver utilization x, which determines waiting times. The relationship between utilization and waiting times has two alternative microfoundations. 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).⁸ 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

$$x = \frac{D(p,x)}{H(p\cdot\theta\cdot x)}.$$
(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 $\frac{\partial H(w)}{\partial w} > 0$, $H(w) \ge 0$, $\frac{\partial D(p,x)}{\partial x} < 0$, and $D(p,x) \ge 0$, 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 depends on both driver

⁸This explanation highlights that the assumption abstracts from scale effects, which refers to the idea that if the number of drivers and passengers in a market doubled, then waiting times would fall (Arnott, 1996; Castillo and Mathur, 2023). I do not have sufficient data to directly test this mechanism, but I examine the estimated model's out-of-sample predictions in subsection 4.4. The model accurately predicts equilibrium responses to an observed policy change, which suggests that scale effects are not strong enough to influence outcomes in policy-relevant counterfactuals. That is, plausible policy interventions are either not sufficiently large to cause meaningful changes in marketplace scale, or scale effects are not very big.

and rider behavior via equation (2). Formally, a platform faces the following problem

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

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,

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

$$-\frac{\theta^*}{1-\theta^*-\tau} + \mu^* \cdot \varepsilon_{D,x} \cdot \varepsilon_{x,\theta} = 0,$$
(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 $\mu = \frac{p \cdot (1 - \theta - \tau) - c}{p \cdot (1 - \theta - \tau)}$.⁹ The latter term is a Lerner-type index (Lerner, 1934), which equals 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 reduce waiting times thanks to the decrease in demand, which spurs an offset in this decline. Equation (5) follows an analogous logic for the setting of commission rates. Raising the commission rate mechanically 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.

⁹For ease of interpretation, I sign all elasticities to ensure that they are positive.

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*

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

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

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

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

Below, I treat $\varepsilon_{D,p}$, $\varepsilon_{D,x}$, and $\varepsilon_{H,w}$ as structural parameters that are invariant to counterfactual scenarios. I discuss the assumption in subsection 2.4, but, given this, two formal definitions are helpful to better understand the implications of theorem 1 and other results below.

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 wages. 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, which is important for pricing. To resolve this, I assume that the platform's make zero profits under perfect competition on both sides of the market, which requires $\varepsilon_{D,p} - \varepsilon_{D,x}$ and $\frac{\varepsilon_{D,p}}{\varepsilon_{D,p}}$ converge to infinity and a constant, respectively.

Passenger prices. The platform sets the price for passengers by marking up their mediation 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. This is reflected by the ratio $\frac{1}{\varepsilon_{D,p}}$ and is socially inefficient.

The optimal price formula is also 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).¹⁰ 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 without reference to drivers. This has important anti-trust implications; abusing labor market power in the gig economy damages both sides of the market, which is an important legal test in the US.11

Commission rates. 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 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 wageposting (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

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

¹⁰The seesaw effect is the tendency for market power over one side of the market to benefit the other side. ¹¹See the US's Supreme Court opinion in Ohio v. American Express Co. (06/25/2018).

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

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

This combines the traditional pricing motivation of a monopolistic firm in a onesided 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.

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 isoelastic assumption for demand and supply, the social planner faces the following problem

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

subject to $D(p, x) = x \cdot H(p \cdot \theta \cdot x).$ (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 of agents on either side of the market governs network effects, precluding scale effects. In the empirical application, this assumption finds support when testing the model out-of-sample, which suggests that policyinduced variation is either insufficient to cause significant scale effects or that scale effects are limited. Therefore, in settings where the ratio of participants can approximate congestion or counterfactual changes in scale are limited, greater competition pushes toward socially efficient outcomes.

More generally, theorem 2 indicates competition may be attractive when it increases the elasticities that platforms face but does not affect the marketplace's scale. This is more likely when market participants can multi-app and where competition comes from the threat of entry by new platforms or from customer adoption of outside options. 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 policies, such as commission caps and minimum wages on utilized hours.

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

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

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

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 as follows

$$\tilde{p} = \frac{1}{1 - \left(\frac{1}{\varepsilon_{D,p}} + \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}\right)} \cdot \frac{c}{1 - \tau} \cdot \frac{\varepsilon_{D,p} - 1}{\varepsilon_{D,p} + \frac{\tau}{1 - \tau}},$$
(16)

$$1 - \tilde{\theta} = 1 - \left(1 - \tau + \frac{1}{\varepsilon_{D,p} - 1}\right) \cdot \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}.$$
(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 already incorporates the social planner's concern that new passengers hurt others through longer waiting times. However, the social planner must correct the monopolistic distortion. Note that because behavioral elasticities are held constant, there is no Spence inefficiency that occurs under the platform's pricing regime.¹²

Second, neither of the socially efficient pricing conditions involves drivers' behavioral responses. This is a consequence of the envelope theorem. Changing either the price or commission rate affects drivers by putting more or less money in their pockets, keeping behavior fixed, and by causing a labor supply response. The former has no effect on the overall surplus because payments in this market are simply a transfer between the platform, passengers, and drivers. Further, workers' behavioral responses have only a second-order effect on their welfare because they are already optimizing.

2.4 Discussion

This subsection discusses several aspects of the model outlined above.

Labor supply. The labor supply function H(w) describes the aggregate number of hours drivers work specifically for the platform in question. In practice, measuring labor supply to a particular platform is difficult because platforms only observe when workers are "online", which measures the hours drivers are logged onto the associated application (Hyman et al., 2020). However, online status is costless to maintain because it does not obligate individuals to do anything. For example, drivers can be at

¹²With heterogeneous demand elasticities, a monopolistic platform would lead to a Spence inefficiency because the platform internalizes network effects for the marginal but not the average rider (Weyl, 2010).

home without the intention of accepting jobs but still appear online, or they may be simultaneously online for a competing platform. Therefore, even in proprietary data, measures of labor supply do not necessarily map to H(w).

This paper exploits the observation that platforms do not need to observe H(w) to optimize. Instead, they can experiment to discover profit-maximizing prices. Therefore, it is possible to infer aspects of platform-specific labor supply by mapping platform choices to, for example, elasticities using insights from the model. Section 4 exploits this intuition to bring the model to the data. The fact that many gig platforms, including Uber, openly experiment with pricing supports the use of this logic in practice.

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 employ a bracketing heuristic whereby they set baseline prices and commission rates to maximize profits and then subsequently finesse their state-dependent pricing.

Costs. In addition to variable costs, platforms may face fixed costs, such as maintaining the code base and data centers that underlie 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, it 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. Fixed costs can also mean that a platform is not profitable, but its behavior is still consistent with the model above.

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). I argue that 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' on-boarding procedures and riders' details are stored on the same server with low marginal costs.¹⁴

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

¹⁴Some empirical evidence of this is the fact that Uber's marketing spend and headcount are only weakly correlated with revenue growth.

Profit maximization. Following much of the literature, the model considers a platform facing a static problem (Jullien et al., 2021). Given the lack of dynamics, it is important to interpret behavioral elasticities as reflecting long-run responses from marketplace participants, which has practical implications in terms of estimation. In particular, it means that elasticities should be informed by responses over a significant time horizon. By virtue of estimating parameters from platform choices, which are informed by long-run elasticities, empirical implementations of this paper's model are immediately reconcilable with profit maximization and suitable for counterfactual policy analysis.

In contrast, existing empirical studies of ridesharing markets generally estimate elasticities to short-run price variation (*e.g.*, surge pricing or week-long pricing experiments) (Castillo, 2023; Rosaia, 2020). However, these short-term elasticities are not the behavioral responses that guide baseline platform pricing. For example, event studies in Hall et al. (2023) illustrate that ridesharing markets take months to adjust. Consequently, the implied elasticities cannot rationalize platform profit maximization, and non-structural assumptions are necessary to evaluate counterfactuals.¹⁵

Constant elasticities. In the social planner's problem and for the counterfactual analysis below, I assume that participants' elasticities are constant. The approach allows for a general interpretation of the results as reflecting changes in equilibrium outcomes under the approximation that these elasticities remain constant. This view is justified by the model's strong out-of-sample performance in section 4, which supports the conjecture that these behavioral responses are relatively stable. However, participants' elasticities will partially reflect strategic interactions with competing platforms. For example, demand may be sensitive to price because price hikes encourage competitors to steal customers by reducing their prices. This makes mapping to welfare quantities via elasticities imperfect (Berger et al., 2022), although it does not affect counterfactual predictions about wages, if treating elasticities as constant is otherwise reasonable.¹⁶

¹⁵These papers assume that platforms maximize a convex combination of profits and participants' welfare and select the weights on the latter to ensure that participant behavior is consistent with platform optimization.

¹⁶One benefit of abstracting from strategic interactions is that it makes the results in this paper comparable with the majority of the literature studying labor market monopsony power, which takes the same approach (Kline, 2025).

3 Redesigning the Marketplace

This section considers alternative market designs to remedy platform monopsony power. Specifically, I consider two policies. First, a cap on commission rates, as has been implemented in other parts of the gig economy, like for restaurants in the food delivery industry (Sullivan, 2022). Second, a minimum wage for utilized hours, which has been enacted by many state and local governments in the US (*e.g.*, most recently in Minneapolis, Minnesota). Lastly, I conclude the section by discussing the relevance of these results and other considerations for policymakers.

3.1 Commission Caps

This subsection studies the introduction of a commission cap into a ridesharing marketplace. Proponents of this policy argue that it offers a way to raise worker welfare. Therefore, I begin by evaluating the optimal commission cap from the drivers' perspective before describing its generic impact on wages and welfare.

Drivers' optimal cap. To determine drivers' preferred commission cap, I consider a three-period game. 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 the same bodies within state and local governments that introduce and enforce existing minimum wages for utilized hours in the gig economy.¹⁷

In *period two*, the platform selects the price for an hour of ridesharing services to maximize its 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 $H(\bullet)$, 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. In this game, the workers must necessarily be better off because the organization can always implement the commission rate that the platform would have wanted to implement.

Backward induction solves the game between the platform and the organization in the following steps. The platform's optimal choice of price given a commission rate

¹⁷In the case of Seattle, Washington, for instance, this would be the City Council.

in equation (4) implies that

$$\varepsilon_{P,\theta} = \frac{\theta}{1 - \theta - \tau}.$$
(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 organization's optimization problem is subject to two constraints. First, utilization will respond to bring the market to equilibrium, which affects wages. Second, the organization 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

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

which yields the first-order condition

$$1 + \varepsilon_{P,\theta} - \varepsilon_{x,p} \cdot \varepsilon_{P,\theta} - \varepsilon_{x,\theta} = 0.$$
(20)

Equation (20) reveals four distinct effects of a commission cap on the wage. First, it mechanically increases wages. Second, it increases wages by encouraging the platform to raise prices. These positive effects on the wage are offset by negative downstream effects. Third, the higher price from the platform's best response reduces utilization and, in turn, wages. Fourth, lower commission rates encourage more supply and decrease utilization and wages further.

Theorem 4 (Wage maximizing commission cap). *The commission rate that maximizes drivers' wages is*

$$1 - \theta^{**} = 1 - (1 - \tau) \cdot \left(\frac{1}{\varepsilon_{D,p}} + \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}\right).$$
(21)

Proof. This follows from plugging equations (6), (7), and (18) into equation (20). \Box

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 optimal cap reduces the commission rate depending on the elasticity of demand to price. The organization knows the platform will increase prices if the commission rate falls. Consequently, when passengers are insensitive to prices, the organization 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 drivers's optimal commission rate, which is described by equation (21), is the social planner's first best commission rate under perfect competition for passengers.*

Proof. Equation (21) converges to equation (17) under perfect competition for passengers. \Box

Theorem 5 has important implications for policymakers seeking to address monopsony power with commission caps when the passenger market is competitive. In this case, it reveals that drivers' and policymakers' goals are aligned. Both groups would want to set the same commission rate. Consequently, policymakers could consult with drivers to determine the optimal cap without concerns that they are trying to manipulate outcomes with strategic responses.

Driver wages and welfare. Translating arbitrary changes in the commission rate to driver wages and, in turn, welfare requires three considerations. First, driver preferences. Under the assumption of an isoelastic labor supply function, drivers' surplus from the ridesharing marketplace equals

$$U(w) = \frac{w \cdot H(w)}{1 + \varepsilon_{H,w}}.$$
(22)

Increases in the wage rate benefit workers directly and via the number of hours worked on the platform. The latter mechanism is especially pertinent in a labor market characterized by free entry. The change in welfare due to an exogenous change in the commission rate (*i.e.*, not because of changes in behavioral elasticities), which is of interest in the counterfactual here, is given by

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

where $\varepsilon_{U,\theta} = \frac{dU}{d\theta} \cdot \frac{\theta}{U}$.

Second, it is necessary to map changes in the commission rate to equilibrium wages. The elasticity of wages to the commission rate is given by

$$\varepsilon_{w,\theta} = 1 + \varepsilon_{P,\theta} - \varepsilon_{x,p} \cdot \varepsilon_{P,\theta} - \varepsilon_{x,\theta}, \tag{24}$$

where $\varepsilon_{w,\theta} = \frac{dw}{d\theta} \cdot \frac{\theta}{w}$. This is simply the left-hand side of the organization's first-order condition in equation (20). 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 because of its equilibrium consequences, and it encourages higher driver supply, reducing utilization.

Third, earnings from the gig economy make up only a fraction of workers' overall income, typically around one-quarter (Anderson et al., 2021). As a result, total driver welfare moves less than one-for-one with variation in the worker surplus from ridesharing. Changes to overall worker welfare can be derived 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). This mapping rests on an assumption of quasi-linear utility and that the elasticity of labor supply to other activities is similar.

3.2 A Minimum Wage on Utilized Hours

This subsection considers setting a minimum wage for workers' utilized hours (MWU). Such policies have been popular among state and local policymakers because, unlike traditional minimum wages, they do not require knowledge of workers' platformspecific labor supply. Instead, it is only necessary to measure the time drivers spend with passengers, which is readily observable. Denoting the level of the MWU 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

$$\varepsilon_{w,\bar{w}} = 1 - \varepsilon_{x,\bar{w}},\tag{25}$$

where $w = p^* \cdot \theta^* \cdot x$, $\varepsilon_{w,\bar{w}} = \frac{dw}{d\bar{w}} \cdot \frac{\bar{w}}{w}$, and $\varepsilon_{x,\bar{w}} = -\frac{dx}{d\bar{w}} \cdot \frac{\bar{w}}{x}$. That is, the MWU mechanically raises drivers' equilibrium wages one-to-one but also leads to an offsetting change in the equilibrium level of utilization. Whether the policy increases wages on net depends on the balance of these forces.

To calculate the magnitude of the offsetting effect, it is necessary to characterize the equilibrium response of utilization. The elasticity of utilization to the MWU (*i.e.*, $\varepsilon_{x,\bar{w}}$) can be expressed in terms of behavioral elasticities after differentiating the equilibrium condition with the MWU substituted in

$$D\left(\frac{\bar{w}}{\theta}, x\right) = x \cdot H(\bar{w} \cdot x), \tag{26}$$

which gives

$$\varepsilon_{x,\bar{w}} = \frac{\varepsilon_{D,p} \cdot (1 - \varepsilon_{\theta,\bar{w}}) + \varepsilon_{H,w}}{1 + \varepsilon_{H,w} + \varepsilon_{D,x}},$$
(27)

where $\varepsilon_{\theta,\bar{w}} = \frac{d\theta}{d\bar{w}} \cdot \frac{\bar{w}}{\theta}$. Equation (27) contains the elasticity of the commission rate to the minimum wage, which describes the composition of the platform's pricing response to the MWU alongside the broader equilibrium adjustments. That is, the elasticity of utilization to the MWU can be rewritten as weighted average of adjustments in utilization to price and commission rate changes $\varepsilon_{x,\bar{w}} = (1 - \varepsilon_{\theta,\bar{w}}) \cdot \varepsilon_{x,p} + \varepsilon_{\theta,\bar{w}} \cdot \varepsilon_{x,\theta}$.

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

$$\max_{p,\theta} \left[p \cdot (1 - \theta - \tau) - c \right] \cdot D(p, x) \text{ subject to } D(p, x) = x \cdot H(p \cdot \theta \cdot x)$$
(28)
and $\bar{w} = p \cdot \theta$.

Platform optimization implies that

$$1 - \theta^{\dagger} = 1 - \frac{\bar{w}}{c + \bar{w}} \cdot (1 - \tau) \cdot \left(1 - \frac{\varepsilon_{D,x} + 1 + \varepsilon_{H,w}}{\varepsilon_{D,p}} \cdot \frac{1}{1 + \varepsilon_{H,w}} \right),$$
(29)

which comes from totally differentiating equation (26) with respect to x and θ .¹⁸ The dagger notation denotes the platform's endogenous choices in this new environment.

Under a MWU, it is still possible to see the impact of monopsony power elevating commission rates via the term $\frac{1}{1+\varepsilon_{H,w}}$. However, this is confounded by another term that includes drivers' platform-specific labor supply. The ratio $\frac{\varepsilon_{D,x}+1+\varepsilon_{H,w}}{\varepsilon_{D,p}}$ measures the strength of demand sensitivity to price against general equilibrium forces. If demand is relatively price-sensitive, then the platform will set a lower commission rate to satisfy the MWU rather than raise prices. This does not mean that a MWU can raise equilibrium wages under these conditions because earnings also depend on passenger prices and utilization. The following theorem addresses this question in the round.

¹⁸Solving for prices readily follows from substituting equation (29) into the minimum wage constraint $p^{\dagger} = \frac{\bar{w}}{\theta^{\dagger}}$.

Theorem 6 (Equilibrium wage increasing MWU.). *The minimum wage on utilized hours leaves workers better off if and only if*

$$c + \bar{w} < \frac{1}{1 - \left(\frac{1}{\varepsilon_{D,p}} + \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}\right)}.$$
(30)

Proof. This follows from setting equation (25) greater than zero, recognizing that $\varepsilon_{\theta,\bar{w}} = \frac{1}{c+\bar{w}}$ from equation (29), and substituting in.

A MWU does not necessarily benefit workers because the platform's pricing response can cause a reduction in utilization that more than offsets the mechanical increase in wages. Whether this is the case depends on the passenger price markup, as shown in the right-hand side of condition (30), being sufficiently large. Intuitively, platforms would rather increase the price than give a larger share of revenue to drivers. Therefore, it is only when customers are insensitive to prices that price increases do not reduce utilization too much, and drivers are left better off.

3.3 Discussion

These results have practical implications for policymakers who are intervening or considering interventions in gig labor markets with the goal of raising worker welfare. With respect to a MWU, these policies are only effective if their mechanical benefit to workers is not outweighed by equilibrium adjustments in utilization. The results above clarify that this is the case when demand is not too sensitive to price. This can be because platforms have monopoly power over consumers and/or passengers are relatively more sensitive to waiting times. Therefore, a policymaker introducing a MWU to increase wages is implicitly expressing a judgment on one of these features of the marketplace.

In contrast, commission caps do not suffer from this tension; they can raise driver welfare regardless of the competitive state of the passenger market. If there is competition for passengers, policymakers can rely on intelligence from drivers to inform the optimal level of a commission cap since their incentives are aligned with the social planner. Moreover, when demand elasticities are very high, the policymaker need not be concerned about the passenger surplus, which necessarily equals zero under perfect competition on the demand side.¹⁹

¹⁹This follows from the fact that the passenger surplus is $\frac{p \cdot D(p,x)}{\varepsilon_{D,p}-1}$ and assumes that demand elasticities reflect preferences rather than strategic considerations on behalf of the platform.

Further, given the power to set commission rates, policymakers can identify the optimal cap. This follows from the fact that if *status quo* prices and commission rates are observed, then experimenting with a commission cap induces sufficient variation to identify participants' behavioral responses. With this knowledge, it is possible to determine the effect of policy interventions following the derivations above. The next section illustrates this logic but exploits random variation in passenger prices instead of driver commission rates.

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. Motivated by the results in section 3, I consider the effect of introducing commission caps and minimum wages on workers. Setting the commission rate to its first-best level increases wages by 14 percent. Conversely, 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.²⁰ Drivers can also work simultaneously for Uber's competitors, like Lyft, giving rise to the aforementioned issues around the measurement of platform-specific labor supply.

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 primarily determined by time and distance. From the passenger's perspective, two components comprised 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, a fixed percentage, was

²⁰A notable exception to this is New York, where the city has implemented a myriad of regulations affecting platform pricing and the onboarding of drivers.

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 at this time (Chandar et al., 2019).

Passenger fares are also affected by Uber's surge pricing, which algorithmically increases base prices when congestion is high on the platform. It is possible to include the equilibrium response of surge pricing into the model, as shown in appendix A; however, this complicates the exposition and does not meaningfully alter empirical results when added into the estimation.²¹ Moreover, statements by Uber's CEO, Dara Khosrowshahi (*e.g.*, "[the platform] optimizes for an average take-rate"),²² as well as the stability of the platform's revenue to gross bookings in public financial filings, supports the birds-eye view that Uber's behavior can be summarized as optimally selecting base prices and commission rates.

Uber, as a global company, became profitable in 2023 after significant cost-cutting and divestments (*e.g.*, from developing its own autonomous vehicles and from its ridesharing business in Singapore). Prior losses due to these endeavors are unconnected to the performance of the firm's ridesharing marketplace in the US, which is responsible for approximately two-thirds of the firm's revenue. Consequently, treating Uber as a profit-maximizing entity in this market is both plausible *ex ante* and predicts the platform's behavior well *ex post*, as shown in the out-of-sample tests in subsection 4.4. Further, the existence of fixed costs does not preclude the possibility that Uber was both operating at a loss and, simultaneously, profit maximizing.

4.2 Data

Three empirical moments are necessary to identify the model's three structural parameters: $\varepsilon_{D,p}$, $\varepsilon_{D,x}$, and $\varepsilon_{H,w}$. Uber's commission rate (*i.e.*, $1 - \theta$) provides the first empirical moment, which relates to these parameters via equation (9). 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 amount 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.

²¹It is possible to empirically incorporate the equilibrium response of surge pricing because Hall et al. (2023) report the elasticity of Uber's surge multiplier to a change in base prices.

²²See Dara's interview with The Rideshare Guy here.

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 2017.²³ 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 of the booking fee from drivers during that period of time,²⁴ 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.²⁵ 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, connecting to participants' behavioral elasticities through equation (8). 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.

²³Some drivers had a lower Uber fee in that year because they were grandfathered in from previous regimes.

²⁴See discussion here.

²⁵The choice of commission rate is also supported by Uber's breakdown of gross bookings in this blog post.

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, which links to the model's parameters through equation (6).²⁶ Hall et al. (2023) report static and dynamic estimates of this statistic, which exploit randomized pricing experiments by the platform. That is, they estimate the *causal* effect of prices on equilibrium utilization. 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 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

$$\varepsilon_{x,p} \approx \hat{\varepsilon}_{x,p} + \frac{\partial \varepsilon_{x,p}}{\partial \varepsilon_{H,w}} \cdot (\varepsilon_{H,w} - \hat{\varepsilon}_{H,w}) = \hat{\hat{\varepsilon}}_{x,p}, \tag{31}$$
where $\frac{\partial \varepsilon_{x,p}}{\partial \varepsilon_{H,w}} = \frac{1 - (\varepsilon_{D,p} - \varepsilon_{D,x})}{(\varepsilon_{D,x} + 1 + \varepsilon_{H,w})^2},$

where $\hat{\varepsilon}_{x,p}$ is the elasticity of utilization with respect to price measured using online hours. So $\hat{\varepsilon}_{x,p}$ 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,²⁷ 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

²⁶In addition, equation (7) characterizes the equilibrium response of utilization to a change in the commission rate and would provide an over-identifying restriction. Unfortunately, I am not aware of any estimates of this statistic. A related moment is the response of utilization to the introduction of tipping in Chandar et al. (2019). It is possible to conceive of tipping as a decrease in the commission rate; a greater portion of the passenger's payment goes to drivers. Consistent with the results below, Chandar et al. (2019) finds almost no response of utilization to a change in the commission rate.

²⁷This utilization rate only includes time with passengers and corresponds to x in the model.

do not offer a particularly negative picture of drivers' earnings.²⁸

4.3 Estimation

I use a generalized method of moments estimator to estimate the model's structural parameters. Precisely, I select $\boldsymbol{\varepsilon} = (\varepsilon_{D,p}, \varepsilon_{D,x}, \varepsilon_{H,w})$ to minimize the distance between $\hat{X} = (\widehat{1 - \theta}, \hat{p}, \hat{\varepsilon}_{x,p})$ and the model's predictions from equations (6), (8), and (9) using the norm $m(\hat{X}, \boldsymbol{\varepsilon})^T \cdot W \cdot m(\hat{X}, \boldsymbol{\varepsilon})$, where

$$m(\hat{X}, \boldsymbol{\varepsilon}) = \begin{pmatrix} \widehat{(1-\theta)} - 1 - (1-\tau) \cdot \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} \cdot \frac{\varepsilon_{H,w}}{1+\varepsilon_{H,w}} \\ \\ \hat{p} - \frac{1}{1 - (\frac{1}{\varepsilon_{D,p}} + \frac{\varepsilon_{D,x}}{\varepsilon_{D,p}})} \cdot \frac{\hat{c}}{1-\hat{\tau}} \\ \\ \\ \hat{\varepsilon}_{x,p} - \frac{\varepsilon_{D,p} + \varepsilon_{H,w}}{\varepsilon_{D,x} + 1 + \varepsilon_{H,w}} \end{pmatrix},$$
(32)

and W is the weighting matrix.²⁹ Although the model is just identified, weighting is helpful because of the finite sample and sign restrictions on elasticities.

I produce standard errors for the estimates by sampling 500 values of $\hat{\varepsilon}_{x,p}$ from a shifted log-normal distribution with a mean of 1.40 and a standard deviation equal to 0.38 (Hall et al., 2023).³⁰ Therefore, these standard errors reflect only statistical uncertainty from the empirical estimate of the elasticity of utilization to price. Still, I refer to this procedure as estimation because it incorporates some statistical uncertainty, although, in practice, it is very close to a calibration exercise. The sensitivity of results to the commission rate and markup is assessed 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. This is unsurprising since the model is exactly identified but not completely trivial because of the finite sample and sign restrictions on the elasticities. Further, other models have struggled to reconcile Uber's behavior with profit maximization (Castillo, 2023; Rosaia, 2020). I attribute this difference to the use of short- versus long-run behavioral responses in estimation.

²⁸See this blog post again.

²⁹I weight the moments according to either the inverse of their estimated variance, in the case of $\hat{\varepsilon}_{p,x}$, or by the inverse of an educated guess at their variance. I derive the latter by assuming that costs (as a share of the fare) and the commission rate have a standard deviation of two percentage points.

³⁰The shift ensures that simulated values do not fall below one, which is necessary to guarantee the platform has an interior solution. This is readily satisfied by the point estimate of 1.40 for $\varepsilon_{x,p}$ from Hall et al. (2023). Further, the standard error around this estimate overstates the true sampling variability because it corresponds to an estimate for the impact at a specific time horizon. For example, averaging the effect over the fifth to sixth month would yield a more precise estimate.

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

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.

Elasticities that exploit variation in surge pricing, or experiments that last a few days, are generally small, implying that Uber has a lot of market power and should charge higher prices. Longer lasting price experiments on the Uber platform have found much larger elasticities (Christensen and Osman, 2023), which are consistent with the results in this paper.

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.³¹ 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.³² In response, Uber raised prices by 40 percent.³³

Interpreting this through the model, it is possible to map Uber's price response to how the policy affected drivers' utilized earnings. Equation (29), 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,

³¹I do not use this event to provide over-identifying restrictions for the estimation because it occurred four years after the other data moments.

³²This 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

³³See this Uber blog post.

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,³⁴ 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,³⁵ 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 outof-sample. Second, Seattle's minimum wage 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.

4.5 Results

Table 2 shows parameter estimates for nine different combinations of Uber's commission rate and costs. This provides greater transparency to the reader, who can decide for themselves the level of uncertainty in commission rates and costs—and the extent to which this uncertainty is correlated. As a guide, the central scenario is highlighted in bold in the middle of the matrix. I find the most likely deviations from this to be the off-diagonal elements. In other words, costs and commission rates are either likely to be positively correlated (*i.e.*, Uber takes more from drivers when they face higher mediation costs) or uncorrelated, but not negatively correlated.

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

³⁴See this Uber blog post

³⁵See this Uber blog post again.

		39%	34%	29%		
_ <u>Costs</u> _	13%	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.90 \ (<0.01)$	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.90 \ (<0.01)$	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.90 \ (<0.01)$		
		$\varepsilon_{H,w} = 2.36 \ (<0.01)$	$\varepsilon_{H,w} = 3.23 \ (0.01)$	$\varepsilon_{H,w} = 4.39 \ (0.01)$		
	18%	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.84 \ (<0.01)$	$rac{arepsilon_{\mathrm{D,x}}}{arepsilon_{\mathrm{D,p}}} = 0.84~(<\!0.01)$	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}} = 0.84 \ (<0.01)$		
		$\varepsilon_{H,w} = 2.92 \ (0.01)$	$arepsilon_{ ext{H,w}}~=~4.27~(0.02)$	$\varepsilon_{H,w} = 6.37 \ (0.03)$		
	23%	$\left \begin{array}{c} \displaystyle \frac{arepsilon_{D,x}}{arepsilon_{D,p}} \ = \ 0.79 \ (<\!0.01) \end{array} ight.$	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}$ = 0.79 (<0.01)	$\frac{\varepsilon_{D,x}}{\varepsilon_{D,p}}$ = 0.79 (<0.01)		
		$\varepsilon_{H,w} = 3.84 \ (0.02)$	$\varepsilon_{H,w} = 6.31 \ (0.03)$	$\varepsilon_{H,w} = 11.63 \ (0.08)$		

Commission rate

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.

rate and rises if Uber is believed to face higher costs. All the estimates imply a considerable degree of monopsony power unless one maintains the implausible assumption that Uber has *both* an unlikely high level of costs and low commission rate.

In contrast, all of the 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.³⁶ 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 monopsony power allowing them to impose markups.

The different levels of competition on either side of the market are striking. A fortunate feature of the estimation strategy is that it does not require a market definition, but thinking about outside options is useful to explain this aspect of the empirical results. In the long run, passengers have many alternatives to ridesharing: public transportation, private transportation (*e.g.*, driving one's own vehicle, walking, cycling, *etc.*), and they can adapt plans. Conversely, there are few alternative work

³⁶When multiplied by $(1 - \tau)$, which is close to one, this mapping is exact.

arrangements that deliver the flexibility of working in rideshare for drivers.

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 wages, which corresponds to the first-best commission rate given the competitive passenger market, and a minimum wage for workers' utilized hours. In doing so, I also present the predicted impact on wages of restoring Uber's market for drivers to perfect competition.

5.1 Commission Caps

Given the analysis in section 3, I consider setting the commission rate equal to its socially efficient level, as defined in equation (17), and allowing the platform to respond with passenger prices. Policymakers can feasibly implement this policy in many ways. For example, given control of the commission rate, they can induce sufficient variation to infer an optimal commission cap, and, if the passenger market is competitive, they can use information from drivers to determine the optimal rate as shown in theorem 5. Moreover, existing pricing regulations in this labor market indicate a commission cap is both legally and politically feasible.

Using equations (23) and (24), table 3 presents estimates of the impact of monopsony power on wages w and the aggregate worker surplus U, where the latter has been scaled down to account for gig work's tendency to be a secondary source of income.³⁷ 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

³⁷I assume earnings from the gig economy comprise one-quarter of an individual's income.

		Commission rate			
		39%	34%	29%	
- <u>Costs</u> -	13%	$\%\Delta w = 34$	$\%\Delta w = 23$	$\%\Delta w = 16$	
		$\%\Delta U = 29$	$\%\Delta U = 24$	$\%\Delta U = 21$	
	18%	$\%\Delta w = 23$	$\%\Delta \mathbf{w}~=~14$	$\%\Delta w = 8$	
		$\%\Delta U = 23$	$\%\Delta U = 18$	$\%\Delta U = 14$	
	23%	$\%\Delta w = 14$	$\%\Delta w = 7$	$\%\Delta w = 2$	
		$\%\Delta U = 17$	$\%\Delta U = 12$	$\%\Delta U = 8$	

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.

equation (24) as follows

$$\%\Delta w = \left[1 + \underbrace{(1 - \varepsilon_{x,p}) \cdot \varepsilon_{P,\theta}}_{(1-1,18) \times 2.17 = 0.40} - \underbrace{\varepsilon_{x,\theta}}_{\sim 0}\right] \cdot \underbrace{\%\Delta\theta}_{0.23} \approx 14.$$
(33)

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 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



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.

evidence suggests that the platform materially depresses wages relative to the counterfactual. However, because these estimates incorporate the attenuating effect of fare and utilization adjustments, they are smaller than other papers that combine short-term variation in driver earnings to estimate supply elasticities with traditional wage-posting models (Caldwell and Oehlsen, 2021).

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 idea by tracing out workers' wages as a function of this parameter. Attaining the wage gains that occur under the commission cap scenario requires almost tripling 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 effect of the policy on overall worker welfare. These estimates rely on stronger assumptions than the predicted wage changes since they require drivers' behavioral elasticities to be fully representative of preferences, and they also rest on assumptions about labor markets outside of the gig economy.³⁸ Nonetheless, the results indicate that a commission cap leads to large welfare gains for drivers.

³⁸In particular, scaling down welfare changes by the share of income that workers derive from ridesharing is appropriate if drivers' utility is quasi-linear and their labor supply elasticity to other labor markets is approximately the same as to ridesharing, which seems reasonable given similar estimates of elasticities to other industries (Lamadon et al., 2022).

A 14 percent increase in wages raises the workers' total surplus by close to 18 percent. Naturally, the larger the wage change estimate, the greater the improvement in worker welfare. Significant welfare gains are due in part to the flexibility of gig work; when wages increase, drivers can increase hours freely to satisfy their intratemporal optimality condition.

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 (30) 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 subsection 4.4.

In summary, a minimum wage on utilized hours 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 regulation. 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 earnings while drivers carry passengers. Ultimately, the policy fails to target the welfare-relevant quantity: equilibrium wages.

6 Conclusion

This paper develops a tractable model of a typical gig labor market: ridesharing. The framework reveals how platforms use monopsony power over drivers to mark up commission rates according to the labor supply elasticity that they face. Consequently, estimates of firm-specific labor supply remain an appropriate way to measure monopsony power in the gig economy. However, the multi-sided nature of these markets complicates the final effect on workers' wages. Bringing the model to the data, I find that the US's largest ridesharing platform, Uber, exerts substantial monopsony power such that wages are 15 percent below the competitive benchmark.

I consider the role of feasible policies in ameliorating monopsony power in these

labor markets. The efficacy of minimum wages on utilized hours, where policymakers mandate minimum payments to workers for time spent with customers, rests crucially on the capacity of the consumer market to absorb price increases. If price increases significantly reduce demand, then these policies actually reduce equilibrium wages. I find that this is likely the case in the US ridesharing market, despite the policy's implementation across many state and local jurisdictions.

Conversely, the ability of commission caps to raise worker welfare is robust to the competitive state of the passenger market, and the model reveals additional appealing features of this policy. For example, given the power to set commission rates, policymakers can induce variation sufficient to identify the optimal cap. Further, when the passenger market is competitive, drivers would collectively set the first-best commission rate if they had the choice. Therefore, since this paper's empirical results indicate that Uber faces strong competition for its customers, policymakers could be reliably informed by intelligence from drivers when setting a commission cap.

References

- Anderson, M., McClain, C., Faverio, M., and Gelles-Watnick, R. (2021). The state of gig work in 2021. *Pew Research Center*, 8.
- Arnott, R. (1996). Taxi travel should be subsidized. *Journal of Urban Economics*, 40(3):316–333.
- Azar, J. and Marinescu, I. (2024). Monopsony power in the labor market: From theory to policy. *Annual Review of Economics*, 16(1):491–518.
- Belleflamme, P. and Toulemonde, E. (2009). Negative intra-group externalities in twosided markets. *International Economic Review*, 50(1):245–272.
- Berger, D., Herkenhoff, K., and Mongey, S. (2022). Labor market power. *American Economic Review*, 112(4):1147–1193.
- Berger, D., Herkenhoff, K., and Mongey, S. (2025). Minimum wages, efficiency, and welfare. *Econometrica*, 93(1):265–301.
- Caldwell, S., Dube, A., and Naidu, S. (2024). Monopsony makes it big. *Journal of Economic Literature*.
- 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.
- Castillo, J. C. and Mathur, S. (2023). Matching and network effects in ride-hailing. In *AEA Papers and Proceedings*, volume 113, pages 244–247. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- 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.

- 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.
- Diamond, R., McQuade, T., and Qian, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, 109(9):3365–3394.
- Dube, A. (2019). Minimum wages and the distribution of family incomes. *American Economic Journal: Applied Economics*, 11(4):268–304.
- 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.
- Glaeser, E. L. and Luttmer, E. F. P. (2003). The misallocation of housing under rent control. *American economic review*, 93(4):1027–1046.
- 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.

- Horton, J. J. (2025). Price floors and employer preferences: Evidence from a minimum wage experiment. *American Economic Review*, 115(1):117–146.
- 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*.
- Karle, H., Peitz, M., and Reisinger, M. (2020). Segmentation versus agglomeration: Competition between platforms with competitive sellers. *Journal of Political Economy*, 128(6):2329–2374.
- Kingman, J. F. (1961). The single server queue in heavy traffic. *Mathematical Proceed*ings of the Cambridge Philosophical Society.
- Kline, P. (2025). Labor market monopsony: Fundamentals and frontiers. *NBER Work-ing Paper*.
- 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.
- Lee, R. S. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review*, 103(7):2960–3000.
- Lerner, A. P. (1934). Economic theory and socialist economy. *The Review of Economic Studies*, 2(1):51–61.
- Li, Z. and Wang, G. (2024). Regulating powerful platforms: Evidence from commission fee caps. *Information Systems Research*.
- 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.
- Manning, A. (2021). Monopsony in labor markets: A review. ILR Review, 74(1):3–26.
- Neumark, D. and Shirley, P. (2022). Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the united states? *Industrial Relations: A Journal of Economy and Society*, 61(4):384–417.
- 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.
- 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. (2007). An empirical analysis of payment card usage. *The Journal of Industrial Economics*, 55(1):1–36.
- 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*.
- Vergara, D. (2023). Minimum wages and optimal redistribution: The role of firm profits. Technical report, Working Paper. Online Appendix A. Analytical results.

- Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.
- Yu, C. (2024). The welfare effects of sponsored product advertising. *Available at SSRN* 4817542.

Appendices

A Model with a Surge Multiplier

I write prices as a multiplicative function of a base price b and an average surge multiplier m. To capture the role of surge multipliers in helping to equate supply and demand, the average multiplier depends on the utilization rate m = M(x). Therefore, riders face a price

$$p = M(x) \cdot b. \tag{34}$$

Platforms select a base price and commission rate while taking the surge pricing algorithm as given. Formally, platforms face the following problem

$$\max_{b,\theta} \left[p \cdot (1 - \theta - \tau) - c \right] \cdot D(p, x) \text{ subject to } D(p, x) = x \cdot H(p \cdot \theta \cdot x)$$
(35)
and $p = M(x) \cdot b$,

where c and τ correspond to mediation costs as in the main text.

Platform optimization yields two first-order conditions b and θ , respectively,

$$(1 - \varepsilon_{M,x} \cdot \varepsilon_{x,b}) + \mu^* \cdot (\varepsilon_{D,x} \cdot \varepsilon_{x,b} - \varepsilon_{D,p} \cdot (1 - \varepsilon_{M,x} \cdot \varepsilon_{x,b})) = 0,$$
(36)

$$-\varepsilon_{M,x} \cdot \varepsilon_{x,\theta} - \frac{\theta^*}{1 - \theta^* - \tau} + \mu^* \cdot (\varepsilon_{D,p} \cdot \varepsilon_{M,x} \cdot \varepsilon_{x,\theta} + \varepsilon_{D,x} \cdot \varepsilon_{x,\theta}) = 0, \quad (37)$$

where $\varepsilon_{M,x} = \frac{\partial M(\bullet)}{\partial x} \cdot \frac{x}{M(\bullet)}$, $\varepsilon_{D,x} = -\frac{\partial D(\bullet)}{\partial x} \cdot \frac{x}{D(\bullet)}$, $\varepsilon_{D,p} = -\frac{\partial D(\bullet)}{\partial b} \cdot \frac{p}{D(\bullet)}$, $\varepsilon_{x,b} = -\frac{dx}{db} \cdot \frac{b}{x}$, $\varepsilon_{x,\theta} = -\frac{dx}{d\theta} \cdot \frac{\theta}{x}$, and $\mu = \frac{p \cdot (1 - \theta - \tau) - c}{p \cdot (1 - \theta - \tau)}$. Comparative statics on the market-clearing condition provide two more equalities that connect all the relevant elasticities

$$\varepsilon_{x,b} = \frac{\varepsilon_{D,p} + \varepsilon_{H,w}}{\varepsilon_{D,x} + 1 + \varepsilon_{H,w} + \varepsilon_{M,x} \cdot (\varepsilon_{D,p} + \varepsilon_{H,w})},$$
(38)

$$\varepsilon_{x,\theta} = \frac{\varepsilon_{H,w}}{\varepsilon_{D,x} + 1 + \varepsilon_{H,w} + \varepsilon_{M,x} \cdot (\varepsilon_{D,p} + \varepsilon_{H,w})},$$
(39)

where $\varepsilon_{H,w} = \frac{\partial H(\bullet)}{\partial w} \cdot \frac{w}{H(\bullet)}$. Combining equations (36) to (39) to produce yields that platform's optimal choice of price and commission rate.