Financing the Gig Economy

Greg Buchak[*]

February 20, 2020

Abstract

Gig economy technologies like Uber rely on consumer finance markets to allocate durable goods to households for use in both consumption and production. Difference-in-difference evidence showing borrowing constraints limit the increases in vehicle sales, loans, utilization, and employment that typically accompany ride-share entry motivates a structural model quantifying how consumer finance access facilitated these technologies' widespread adoption. Low-cost financing allowed low-income households to obtain vehicles for gig economy employment, and gig economy entry increased welfare by $3 billion annually when compared to traditional forms of production. Historically representative financing costs would have prevented low-income households' entry and eliminated these gains.

[*]Stanford University Graduate School of Business. (buchak@stanford.edu). I am grateful to Amit Seru, Amir Sufi, Gregor Matvos, Steve Kaplan, Ufuk Akcigit, Lisa Bernstein, Douglas Baird, Tomasz Piskorski, Zhiguo He, Luke Taylor, Lars Hansen, Eric Zwick, and participants at the NYU Salomon Center/BPI Conference on Household Finance, the Federal Reserve Biennial Conference on Auto Lending. Parts of this paper were made possible with a research grant from the Fama Miller Center at the University of Chicago Booth School of Business. This paper uses data that is calculated (or derived) based on credit data provided by TransUnion, a global information solutions company, through a relationship with the Kilts Center for Marketing at the University of Chicago Booth School of Business.

1 Introduction

Firms like Uber, Lyft, and Airbnb have created convenient markets for households to sell capital services with their durable consumption goods. Car owners can provide rides, and homeowners can house guests in their spare rooms. As the reach of these technologies has expanded, individuals have joined the so-called gig economy in droves as drivers, delivery people, and hosts. As of 2019, Uber estimates that roughly three million Uber produced roughly $30 billion in North American gross bookings, while Airbnb hosts listed more than 7 million vacancies.[1] While there has been intense debate on the role of technology and regulatory arbitrage in the gig economy's growth, this paper focuses on a key distinguishing characteristic of gig-economy firms: in contrast to traditional forms of production, gig economy workers must bring their own—often costly—physical capital to work.[2]

Organizing production around workers supplying capital is a blessing and a curse. Gig economy production allows households to extract capital income from their durable consumption goods, but because gig economy employment pays a low wage, would-be drivers tend to come from the lower end of the income and skills distribution.[3] Consequently, they often lack cars and must obtain them through costly borrowing. This paper quantifies the welfare and distributional consequences of this trade-off. Given the low-cost financing that was available when gig economy technologies arose, the dual asset use benefit dominated. Low-income households were able to obtain capital and gig economy employment, and welfare increased by $3 billion relative to traditional forms of production. However, had financing costs been at historical averages, the welfare gains from gig economy production compared to traditional firms would have been zero, and low-income, and finance reliant households would have been excluded from participating. Hence, while the

[1] See `https://www.theguardian.com/us-news/2019/mar/22/uber-lyft-ipo-drivers-unionize-low-pay-expenses` and `https://news.airbnb.com/fast-facts/`, visited December, 2019.

[2] In the case of Uber or Lyft, they must *own* the car—leasing or renting is not sufficient. Standard leases and rental agreements have mileage caps and rules against commercial use. Uber attempted to provide leases itself, although the program infamously lost money and was closed.

[3] See Chen et al. (2019). Also Mishel (2018) estimating the average after-expense Uber income is $11.77 per hour.

gig economy's organization of production offers large potential benefits through dual asset use, particularly to low-income households, these benefits rely critically on households' borrowing ability.

I first highlight a mismatch in ex-ante capital ownership leading to a reallocation after ride share entry. In a difference-in-difference analysis exploiting the differential timing of ride share entry across cities in the United States between 2010 and 2016, I show that entry leads to a 1.6% increase in new auto sales, a 0.60% increase in employment, and an additional 140 miles driven per day per registered ride share driver. These effects are large relative to the total number of ride share drivers, which is roughly 0.9% of the employed population in mature markets. These results are largest in low-income zip codes, consistent with ride share driving being most attractive to low-income individuals. Additionally, the increases occur only among cars meeting the specific physical criteria set forth by Uber and Lyft.

Crucially, this growth relies on financing. The same difference-in-difference approach shows that auto loan originations increase by 1%, a magnitude that implies that the majority of new auto sales are financed. These increases are concentrated in low-income zip codes and among ride-share eligible vehicles. To show that financial constraints mute the impact of ride share entry, I exploit a provision in the Fair Credit Reporting Act requiring credit agencies to remove bankruptcy history 10 years after filing. Consumers who recently had their bankruptcy history purged follow the same difference-in-difference response to ride share entry, while those who have not yet had the information removed do not respond at all. At the zip-code level, areas that historically relied on banks for auto credit but saw credit contractions following post-crisis regulatory changes[4] see smaller responses in terms of auto loans, sales, and employment growth. Similarly, areas with large mortgage delinquency rates in the crisis that consequently have impaired credit when ride share enters see reduced real effects.

These reduced form facts evidence a significant reallocation of capital towards low income drivers relying on consumer finance. However, the analysis is silent about the costs and benefits of this form of production in comparison to counterfactual methods of organization and the equilibrium effects of access to finance. To

[4]See Benmelech et al. (2017) and Buchak et al. (2018a).

quantify these effects, I build a structural model in the spirit of Berry et al. (1995) more recently applied in a consumer finance context by Egan et al. (2017), Benetton (2018), and Di Maggio et al. (2019). I extend these models to explicitly capture the consumer financing and ride share labor market decisions that lie at the heart of my analysis. I further augment the traditional estimation approach by using rich micro data on vehicle utilization, financing, and ride share driving. I close the model with a simple model of ride demand and study the equilibrium outcomes arising in these interconnected markets under counterfactual organizational structures and financing environments.

The model suggests that compared to an industry structure where firms own the capital but have otherwise identical technology and regulation, the gig economy mechanism leads to an increase in consumer welfare of $3 billion by offering 11% more rides at 26% lower prices. Annual vehicle utilization increases by roughly 15 billion miles, relative to roughly 3 trillion total in the United States. This occurs because in contrast to traditional forms of production, gig economy capital provides both consumption benefits and capital income, and hence payment for capital services alone need not justify vehicle purchases. Consistent with empirical evidence showing that high-income households do not reduce car purchases after Uber enters despite using them less, the model shows that the majority of the consumption benefit of car ownership obtains from the convenience of owning the car, and is relatively insensitive to actual utilization. Hence, using the car for ride share does not reduce the consumption benefit of car ownership, and car owners can enjoy both benefits simultaneously.

These benefits come at the cost of a reliance on low-income consumers' access to finance. Ride share arose in an era of cheap auto finance, and had ride share arisen when the costs of consumer finance were more in line with historical averages, its method of capital allocation would have been roughly equivalent to a more traditional method in terms of prices and quantities. Financing costs additionally drive important wage and demographic dynamics. Counter-intuitively, as financing costs rise, driver net income rises even accounting for financing costs, because low-income, finance reliant drivers are forced out of the market. Higher-income, less finance-reliant individuals enter but demand higher wages. Hence, while wages rise, the low-income drivers who exited the market do not benefit. Removing access to consumer finance entirely reduces welfare by more than $3 billion but doubles driver net income for those drivers able to obtain cars without financing.

These results highlight a key drawback of gig economy production in that it requires drivers to own their cars. I therefore consider the impact of emerging technologies that introduce short-term, on-demand rental markets that connect car-owning households with drivers that alleviate this ownership-use friction.[5] The successful introduction of such a technology would lead to a $7 billion welfare gain in the ride share industry. This improvement arises because high income, low-cost-of-capital households own cars and derive large durable consumption benefits without facing high finance costs, while low-cost-of-time workers rent and drive them, unencumbered by the burden of capital ownership. Importantly, this counterfactual assumes there are no transaction costs or other frictions such as a moral hazard between the driver and owner. Because these frictions almost surely exist in practice (and likely explain why Uber's rental business failed), an alternate interpretation is that this quantification places a lower bound on the size of these frictions.

More generally, my model highlights and quantifies important costs and benefits in reorganizing production so that capital owners use capital for production and for durable consumption. Beyond ride share, the findings in this paper help explain why other technologies that allow durable consumption goods to be used to produce capital income—for example, AirBnB—have succeeded, while other seemingly similar technologies that do not—for example, WeWork—have not. Finally, my paper highlights an important and novel role that consumer finance plays in directly facilitating growth by enabling household production and the widespread adoption of disruptive technological innovations.

1.1 Related literature

My paper ties together questions and methodologies from household and corporate finance. Like recent work in household finance (e.g., Benetton et al. (2018), Benetton (2018), Egan et al. (2017), Di Maggio et al. (2019), and Buchak et al.

[5] Apps, such as Turo, `https://turo.com/`, are starting to offer these services.

(2018b)) I use Industrial Organization demand estimation techniques to study consumer behavior in financial markets. In contrast to these papers, I focus on a consumer financing a capital good for production.[6] My paper extends traditional approaches by incorporating an explicit financing and labor market choice, allowing me to study the joint relationship between financing, employment, and productivity.

More generally, studies of consumer finance typically center on how the financial system and financial constraints impact household consumption, durable goods purchases, and savings, e.g., in Campbell (2006). This arises in the context of cars, e.g., Grunewald et al. (2019), Benmelech et al. (2017), housing, Mian and Sufi (2011), Piskorski and Seru (2018), Hurst and Stafford (2004), Di Maggio et al. (2017), consumer loans, Zinman (2010), Melzer (2011). My paper studies similar financial decisions but for the purpose of production. Similarly, recent studies have examined the connection between employment and consumer finance, e.g., in Dobbie et al. (2016), Donaldson et al. (2018), and Herkenhoff et al. (2016). Unlike these papers, my study examines the direct link between household financial access and their ability to work in a market that requires household capital. This finding complements literature evaluating the size and benefits of the financial system, such as Greenwood and Scharfstein (2013) and Philippon (2015), by showing another important role of consumer finance.

My paper additionally speaks to a long line of research emphasizing the importance of the financial system in growth, e.g., as in Mian et al. (2017), Kaplan and Zingales (1997), Jayaratne and Strahan (1996), the role of finance in productivity, i.e., in remedying the sorts of misallocation as in Hsieh and Klenow (2009), such as in Lenzu and Manaresi (2017), Buera et al. (2011), or Midrigan and Xu (2014). My paper touches on these issues by highlighting the importance of ex-ante factor misallocation at the time of the introduction of a disruptive technological change that impacts how capital can be used. Here, the consumer finance system takes on a role typically played by corporate financial markets to reallocate factors of production towards more efficient users.

Finally, many empirical papers have studied the impact of ride share on con-

[6]Indeed these models have their roots in the models of Berry et al. (1995) and Nevo (2001) which model consumer vehicle purchasing decisions for durable consumption.

sumer surplus, as in Cohen et al. (2016) and Cramer and Krueger (2016), on labor markets, as in Hall et al. (2017), Benjaafar et al. (2018), Cook et al. (2018), Cook et al. (2019), and Chen et al. (2019), or on traffic and vehicle crashes, as in Barrios et al. (2018). In contrast, my paper focuses on the impact and implications of the capital stock, with a particular focus on the benefits of dual asset use and consumer finance. Other papers have explored the theory behind this question. Horton and Zeckhauser (2016), Fraiberger and Sundararajan (2017), Razeghian and Weber (2016), and Ostrovsky and Schwarz (2018) consider theoretical models of vehicle ownership and utilization in a ride share economy. I evaluate empirically the typical prediction that ride share entry will lead to a smaller capital stock through greater utilization rates. I show empirically that this prediction is complicated by the large option value that households obtain from car ownership beyond utilization as well as the ex-ante misallocation of capital.

2 Data and institutional background

Data: This section describes the data, and Appendix 6.1 provides greater detail and summary statistics. I obtain the staggered entry dates of Uber and Lyft by hand-collecting company press releases and newspaper articles. I use a proprietary dataset from Uber providing the number of registered drivers at a CBSA-month level. I obtain data on auto sales, registrations, and auto loans from RL Polk, the North Carolina, Washington, and Indiana DMVs, and Equifax, respectively. RL Polk and Equifax are nationwide datasets providing the number of new vehicle sales and auto loans at the zip code level between 2010 and 2016. The DMV data are vehicle-level registration data by zip and month, which I join to data from the National Highway Traffic Safety Administration and FuelEconomy.gov databases to obtain the physical attributes of the car, including make, model, year, horsepower, and fuel efficiency. I use borrower-level data from a 10% sample of all borrowers in TransUnion, which provides the individual's borrowing activity and information on her past bankruptcy filings. Finally, I use public government databases, including the IRS Summary of Income, and the United States Census and American Community Survey (ACS).

Institutional background: Uber began operations in San Francisco in 2010, with Lyft following shortly thereafter. Both services expanded rapidly to other cities. By the end 2016, there were nearly 800,000 registered Uber drivers.[7] Ride share's entry is not random, which presents an identification challenge if entry is correlated with local conditions in industries related to auto sales or loans. Appendix Section 6.2 provides details on the timing and determinants of ride share's entry. In short, entry is more likely in large cities with high mobile broadband penetration, suggesting that these services entered areas with large potential markets. Vehicle ownership rates or access to finance do not predict entry. These findings hold both on the extensive margin of entry and in terms of the timing of entry.

3 Ride share's real effects and the role of finance

This section studies the empirical effects of ride share entry and finance's role in these effects. I first study ride share's impact on sales, loans, employment, and vehicle utilization before examining how financial constraints inhibit these effects. As a threshold matter, it is important to note that while the overall rate of vehicle ownership in the United States is high, ownership rates are heterogeneous across the income distribution, with households earning below the full-time equivalent Uber income owning roughly 0.20 fewer cars per adult household member than households earning at the median income. Appendix Section 6.3 details these facts.

3.1 What happens when ride share enters?

I begin by analyzing whether ride share's entry prompted lower-income households to obtain cars, how they financed them, and whether entry corresponded to increases in vehicle utilization and employment. These analyses share a common difference in difference empirical design that exploits ride share's staggered entry across MSAs at the zip-time level. All zip codes in the sample eventually see ride share entry; identifying variation comes from comparing zip codes that ride share has already entered to zip codes that ride share has not yet entered. I further exploit

[7] This number does not include drivers registered for Lyft but not Uber. Mishel (2018) estimates that the entire sharing economy is roughly 50% larger than Uber alone.

within-city variation in median incomes for a third difference, examining whether the impact of treatment is concentrated in relatively low-income zip codes. The main specifications are as follows:

$$Y_{zt} = \beta Post_{zt} + \gamma_t + \gamma_z + \epsilon_{zt} \tag{1}$$

$$Y_{zt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Low\ Income_z + \gamma_z + \gamma_{Income,t} + \epsilon_{zt} \tag{2}$$

Where $Y_z t$ is the outcome variable of interest at zip z and time t. $Post_{zt}$ is an indicator for ride share entry, $Low\ Income_z$ is an indicator for whether the zip code's median income is in the bottom 50% of zip codes in the MSA, γ_z are zip fixed effects, γ_t and $\gamma_{Income,t}$ are quarter fixed effects and quarter $\times$ $Low\ Income_z$ fixed effects, respectively. I also present event studies that replace $Post_{zt}$ with an set of event time indicator variables. These event studies are helpful in addressing the primary identification concern, which is the endogenous timing of ride share entry by showing pre-trends. The results follow.

3.1.1 Auto sales and loans

I begin by studying the impact of ride share entry on auto sales and loans. The outcome variable is log new auto sales from RL Polk, and log new auto loan originations from Equifax. Table 1 Panel A shows the results for new auto sales. Column (1) shows a 1.6% increase in vehicle sales following ride share entry. Column (2) shows the within-CBSA regression, and finds that following ride share entry, low-income zip codes see auto sales increase by roughly 2.6%, while high-income zip codes see auto sales increase by 0.6%. These findings are robust to the inclusion of time fixed effects that vary with income group, showing that the results are not driven by time trends that vary between low- and high-income zip codes unrelated to ride share entry. Panel B shows the results for new auto loans. The results mirror those for new auto sales, indicating that the increases in physical car sales are accompanied by significant increases in vehicle financing.

Figure 1 presents the event studies. Panels (a) and (b) show the results for new auto sales and new auto loans for all zip codes, respectively. These figures show no pre-trends with an increase in auto sales and loans that occur around the time of entry and follow similar dynamics to the entry of drivers shown in Figure 3 Panel (c). Panels (c) and (d) show these results among low-income zip codes, and find stronger effects, while Panels (e) and (f) show the results among high-income zip codes, and find no effects, as expected.

As a second robustness check, I first exploit ride share vehicle eligibility requirements and detailed vehicle-level registration data. To be eligible for ride share, a vehicle must be no older than 15 years, have four doors, and be a sedan, SUV, or minivan. In Appendix Section 6.4, I show that increases in vehicle registrations are confined entirely to eligible vehicles in low income zip codes. Additionally, I observe whether the registered vehicle has a lien, indicating that it is financed, and I find that the probability of having a lien increases following ride share entry, but only for ride share eligible vehicles in low-income zip codes.

As a third robustness check, I repeat the preceding analysis but split zip codes on 2010 transportation worker share rather than wage to confirm that the zip-level income measure captures variation relative to ride share *driving* in particular. The results, shown in Column (3) of Table 1 Panels A and B, show that auto sales and loan increases are concentrated in zip codes with a high share of transportation workers. As a final robustness check, I run a placebo test of (1) and (2) where the dates of ride share entry are randomly assigned among to different locations. The results in Appendix 6.10 show no significant effect, confirming that the preceding findings are related to ride share entry and not uncontrolled-for time trends in the data.

3.1.2 Employment

I next test whether ride share entry coincides with increases in low-income employment by using data from the IRS Summary of Income. These data show the number of tax filings in the zip code and year, broken into by adjusted gross income (AGI) buckets. The specification follows (1) and (2) but uses log number of tax filings as the outcome variable,[8] is at the zip-AGI bucket-year level, includes

[8]By using the number of tax filings as the outcome variable, I am measuring employment impacts on the extensive margin. Note that these impacts may also include transition from the untaxed sector of the economy into the taxed sector of the economy.

year-AGI and zip-AGI fixed effects, and the low-income differencing variable is an indicator for whether the AGI bucket is below $25,000 per year rather than a differencing variable at the zip-code level. $25,000 is the relevant split because full-time ride share driving pays slightly below this level on average.

Table 2 Panel A Column (1) shows an increase in tax filings across all AGI brackets of roughly 0.6% following ride share entry. Column (2), which includes the low-income AGI interaction, shows that the increase in filings is concentrated entirely in filings in the $\leq \$25,000$ AGI bucket, leading to an employment increase of $1.1%. Column (3) includes zip times year fixed effects to absorb any time-zip varying economic conditions correlated with the endogenous entry choice and shows that the differential effect between low- and high-AGI filings persists.

To check for pre-trends, Figure 2 Panel (a) shows the results of the event study for all filings aggregated at the zip-year level across buckets, together with with 95% confidence intervals. The figure shows no pre-trends. Once ride share enters, tax filings increase. Panels (b) and (c) split the sample into filings below $25,000 and filings above $25,000, respectively and confirm that tax filings increase for low-income filers but not for higher-income filers, with no pre-trends. As a second robustness check, I utilize the fact that ride share drivers report earnings as business, rather wage income, which the IRS data reports separately, and Appendix 6.6, confirms that the increases are driven by increases in business tax filings. Finally, placebo tests reported in Appendix Section 6.10 shows no effect.

3.1.3 Vehicle utilization

I next test whether these ride share cars see higher utilization rates after entry. The South Carolina DMV data report vehicle odometer readings when vehicles ownership. I calculate yearly utilization rates for completed ownership spells and test whether utilization rates of ride-share eligible cars change following ride share entry in low-income zip codes. The specification again follows (1) and (2) but is run at the vehicle-ownership spell-zip level with interactions and fixed effects described below.

Table 2 Panel B shows the results. Column (1), which shows the overall effect of ride share entry, contains zip-eligibility and quarter-eligibility fixed effects.

The effect is positive but insignificant, and implies that the typical ride share driver drives roughly 140 miles per day.[9] Columns (2)-(6) include more interactions and fixed effects. Focusing on Column (5), which includes the triple interaction of entry, the low-income zip indicator, and the vehicle eligibility indicator, the estimates show a large and statistically significant impact of ride share entry on vehicle utilization in low-income zip codes among ride share-eligible cars of roughly 2,000 miles per year. Placebo tests reported in Appendix Section 6.10 show no effect.

These positive changes are offset in aggregate by large decreases in utilization among ineligible vehicles and high-income zip codes. Importantly, while vehicle utilization *decreases* on average in high-income zip codes, earlier results show that vehicle sales in high-income zip codes do not decrease, suggesting that vehicle utilization may not be an important factor in a consumer's decision to purchase a car for durable consumption. I return to this finding later.

3.2 Financial constraints and ride share growth

The previous section documented the effects of ride-share entry and showed that auto sales, loans, employment, and utilization increase significantly. The large contemporaneous rise in auto loans suggest that a lack of credit may dampen these effects. To study this question, I exploit the Fair Credit Reporting Act's (FCRA) requirement that credit agencies remove Chapter 7 bankruptcy filings from credit reports ten years after filing. This policy generates exogenous variation in borrowing costs between individuals who have or have not yet had this bankruptcy information removed, and I test whether those individuals with higher borrowing costs borrow less in response to ride share entry and find that they do. I then bring these findings to the zip-code level to test whether variation in credit access leads to smaller real effects. I generate zip-level credit access variation using (1) the share of consumer loans that became seriously delinquent in 2010, and (2) the 2010 bank share of auto lending. Indeed, these measures predict less auto loan growth and lead to smaller real effects.

[9] 175 extra miles per year, divided by the 0.005 overall share of ride share drivers, divided by 252 working days gives 140 miles per day per ride share driver.

3.2.1 Borrower-level evidence from bankruptcy flag removal

I begin by exploiting the exogenous removal of bankruptcy filing information from consumer credit records. The FCRA requires that credit bureaus remove Chapter 7 bankruptcy filing reports no more than ten years after filing. Previous literature, e.g., Musto (2004) and more recently Herkenhoff et al. (2016) and Dobbie et al. (2016), has documented that this removal leads to a large and statistically significant increase in credit scores and a decrease in borrowing costs. I first confirm these findings in the context of auto loans specifically. As motivating evidence, Figure 3 Panel (a) shows the probability of receiving an auto loan in a given quarter versus the number of years since the Chapter 7 filing. There is a large discontinuity in the probability of obtaining an auto loan precisely ten years following the filing, consistent with the predictions and previous literature. To adjust for time trends in the rate of auto lending, I run the following regression to isolate the impact of filing removal:

$$Origination_{izt} = \beta I(YearsSinceFiling \geq 10) + \gamma_{zt} + \epsilon_{izt} \tag{3}$$

Table 3 Panel A shows the results for regression (3). Columns (1)–(4) vary the window around the ten year cutoff from ± 0.25 years to ± 2.50 years. These results consistently show a 10-15 basis point increase in the quarterly probability of obtaining an auto loan after flag removal. Figure 3 Panel B shows these results graphically, replacing the post-removal dummy with a series of event-time dummies and plots the coefficients. The result shows a clear and statistically significant discontinuity with no systematic trends before or after the cutoff.

I next examine how these differences translate into differential borrowing responses to ride share entry. I form two groups of borrowers. The *constrained* group of borrowers filed for bankruptcy between 8 and 9 years prior to ride share entry, 9 not included. The *unconstrained* group of borrowers filed for bankruptcy between 11 and 12 years prior to ride share entry, 11 not included. I study the borrowers' auto loan originations in an event window one year before and after ride share entry. This timing implies that borrowers in the constrained group had the bankruptcy flag for the entire event window, and that borrowers in the unconstrained group did not have the bankruptcy flag at any time during the event window. Otherwise, borrowers in these groups are as close as possible to each other in terms of when they filed for bankruptcy. Appendix Section 6.7 details the timing of this experiment. I run the following regression:

$$Origination_{izt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Constrained_i + \gamma_{gt} + \gamma_{gz} + \gamma_{zt} + \epsilon_{izt} \quad (4)$$

The regression is at the borrower-quarter level. $Origination_{izt}$ is an indicator for whether the borrower obtains an auto loan in the quarter. $Post_{zt}$ is an indicator for whether ride share has entered zip z as of time t. $Constrained_i$ is an indicator for whether the borrower is in the constrained group. γ_{gt} is a time-constraint group fixed effect; γ_{zt} is a zip-constraint group fixed effect; γ_{zt} is a time-zip fixed effect. All borrowers in the regression fall in either the constrained or unconstrained group and therefore, all declared bankruptcy at most four years apart and at least eight years prior to the event window over which they are being compared.

Table 3 Panel (B) shows the results, with coefficients shown in percentage terms. Column (1) is the overall level effect which, while positive, is insignificant for this particular subsample of borrowers.[10] Column (2) includes the interaction of treatment with whether the bankruptcy flag is visible, and shows a statistically significant impact of ride share entry of roughly 24 basis points for the unconstrained control group and a close-to-zero effect for the constrained treatment group. The differential effect survives the inclusion of zip-time fixed effects in Column (3). For robustness, Columns (4)–(6) repeat the same exercise with a narrower window, where constrained borrowers declared bankruptcy between 8.5 to 9 years prior to ride share entry and unconstrained borrowers declared bankruptcy between 11 to 11.5 years prior to ride share entry. The estimates are unchanged.

Figure 3 Panels (c) and (d) show the event study, calculated as the accumulated difference in difference coefficients over the event window for the constrained and unconstrained groups, respectively. These figures show little change in origination probability for the constrained group, while the unconstrained group sees an increase in auto loans beginning on impact. The pre-trends in both groups show no

[10]This smaller effect is itself consistent with the idea that borrowers with worse credit—those that declared bankruptcy within 12 years—respond less to ride share entry than the broader population of borrowers who have not declared bankruptcy.

significant effects leading up to entry. This analysis shows that borrowers facing greater credit costs increase borrowing less in response to ride share entry.

3.2.2 Zip-level evidence

I conclude the analysis by using zip-level variation in credit access to test whether financial constraints reduce the real effects of ride share entry. I obtain variation from two sources: The first comes from the percentage of consumer loans that were in serious delinquency, defined as being 60 days or greater behind on payments, as of 2010. Consumers who were seriously delinquent on 2010 consumer loans will have impaired credit scores and therefore face higher borrowing costs when ride share enters. The second source of variation in credit availability comes from the 2010 bank share of auto lending. Following the financial crisis, new banking regulations such as increased capital requirements and stricter supervision reduced banks' ability to lend.[11] Those zip codes where banks dominated auto lending in 2010 saw their lenders hampered by regulations and face a relatively constrained supply of auto credit.[12]

Following the empirical strategy outlined above, I find that neither zip-codes with a high share of delinquencies nor zip-codes with a high share of bank credit supply see effects of ride share entry in terms of auto loans, sales, and employment growth, demonstrating that access to credit is key for the real effects of ride share entry to take hold. The analysis shows that these zip codes demonstrate no pre-trends, and placebo tests confirm that these differences are driven by the precise timing of ride-share entry. These results are shown and discussed in detail in Appendix Section 6.7. To summarize, the results in this section establish a clear link between the supply of finance and ride-share uptake: low-income households require financing to obtain cars to drive in the gig economy. When they cannot obtain it, gig economy growth and employment is reduced.

[11] For example, the closure of the Office of Thrift Supervision and less favorable capital treatment of mortgage servicing rights. See, e.g., Buchak et al. (2018a).

[12] See Benmelech et al. (2017), which uses *pre*-crisis share of bank lending and finds that higher bank shares lead to more lending ex-post, as non-bank lenders were disproportionately hurt during the crisis. My experiment uses *post*-crisis share of bank lending, the idea being that post-crisis banking regulations disproportionately hurt banks.

4 A structural model of ride share capital

Ride entry permits households to use their durable consumption goods to earn capital income. However, for households to obtain the necessary capital, they often need to borrow, and lack of household financial access becomes an impediment to gig economy participation. In this section, I build a structural model to quantify these costs and benefits and the equilibrium impacts of financial costs on gig economy prices, quantities, and participation across the income distribution. The model builds on Berry et al. (1995) by including a financing and ride share labor supply choice, along with a market for ride share services.

4.1 Model description

Consumer $i \in \{1, \cdots, I\}$ in market $c \in \{1, \cdots, C\}$ makes three decisions: (1) which (if any) vehicle $j \in \{0, 1, \ldots, J\}$ to acquire; (2) whether to finance the purchase or pay cash outright; (3) whether to become a ride-share driver. Utility over each alternative is parameterized by consumer-specific characteristics summarized in θ_i. These characteristics include, for example, her preference over vehicle fuel economy, her available liquidity and financing costs, and her non-monetary preferences over ride share driving. These characteristics are themselves driven by the consumer's underlying demographics, D_i, and the mapping between D_i and θ_i is to be estimated from the data.

A market c is at the zip-quarter level and is summarized by a distribution of consumer demographics, $F_c(D_i)$, and $\phi_c \in 0, 1$, an indicator for ride share presence. Each vehicle $j \in \{0, 1, \cdots, J\}$, with $j = 0$ being the outside option of not owning a vehicle, has observable attributes x_j which include, for example, price and gas mileage. A particular element of x_j is $e_j \in 0, 1$, an indicator for the vehicle's ride share eligibility. With sufficient liquidity to cover the up-front price of the vehicle, the consumer can either pay with cash or finance the vehicle; with insufficiently liquidity, the consumer must finance the vehicle. I close the model with a simple model of ride share demand using elasticities from Cohen et al. (2016).

4.1.1 Individual consumer's problem

The consumer obtains utility from: (1) owning the car for her own durable consumption, (2) disutility from purchasing and financing the vehicle, and (3) utility from ride-share driving. I directly model indirect utility as described below.

Utility from own consumption: Consumer i obtains utility from owning the vehicle for her own durable consumption. Her utility is a function of the vehicle's characteristics and her utilization for her own (non-ride share) purposes. Indirect utility takes the following form:

$$u_c(\theta_i, x_j', \phi_c, \xi_{jc}, \epsilon_{ijc}^c) = \underbrace{\beta_i^0}_{\text{outside option}} + \underbrace{x_j'\beta_i + \xi_{jc}}_{\text{vehicle characteristics}} + \underbrace{\beta_i^m m_i}_{\text{utilization}} + \underbrace{\epsilon_{ijc}^c}_{\text{preference shock}} \tag{5}$$

β_i^0 captures the value of car ownership versus the outside option of not owning a car, common across all vehicles, but allowed to vary before and after ride share entry. The vehicle characteristics term captures utility from car specific characteristics, with x_j' containing observable characteristics such as horsepower, and ξ_{jc} containing characteristics common to the vehicle-market that are not observed by the econometrician, such as quality. β_i is a vector of consumer preferences. The utilization term captures utility from using the car, where m_i is the consumer's desired miles driven and β_i^m is the consumer's utility per mile driven, which can vary before and after ride share entry. ϵ_{ijc}^c is an idiosyncratic shock to consumer i's preference for vehicle j, interpreted as horizontal product differentiation.

Utility from financing: Vehicle j costs p_j, and regardless of her financing decision, the consumer receives disutility $\alpha_i p_j$ from acquiring the car, where $\alpha_i > 0$ is her price sensitivity. Consumer i can buy a car with cash or through financing, and the indirect disutility of obtaining financing has the following functional form:

$$f(\theta_i, p_j, \epsilon_{ijc}^f) = f_i^0 + \epsilon_{ijc}^f \tag{6}$$

f_i^0 captures systematic costs in obtaining financing, for example, financing charges in excess of the consumer's discount rate, origination costs, and search costs. f_i^0 is consumer-specific and so accommodates, for example, high-income borrowers having access to lower-cost credit than low-income borrowers. ϵ_{ijc}^f is a consumer-specific shock to financing utility, which captures, consumer's idiosyncratic financing needs, sophistication, debt aversion, or costs. I assume that being a ride share driver does not impact financing costs.[13]

Conditional on selecting vehicle j, the consumer must obtain financing if her liquidity l_i is insufficient to cover the vehicle's price outright, i.e., if $l_i < p_j$. Otherwise, she has the option to obtain financing or pay outright. Her utility from obtaining and financing a car is summarized below:

$$u_f(\theta_i, x_j', \epsilon_{ijc}^f) = -\alpha_i p_j - \begin{cases} f(\theta_i, p_j, \epsilon_{ijc}^f) & l_i < p_j \text{ (constrained)} \\ \begin{cases} f(\theta_i, p_j, \epsilon_{ijc}^f) & \text{financed} \\ 0 & \text{not financed} \end{cases} & l_i \geq p_j \text{ (unconstrained)} \end{cases} \tag{7}$$

Utility from ride-share driving: Consumer i can become a ride share driver if (1) she lives in a ride share market, $\phi_c = 1$, and (2) the vehicle she purchases is eligible for ride share driving, $e_j = 1$. Her indirect utility from driving or not driving is:

$$u_{rs}(\mathbf{p}, \theta_i, x_j', \phi_c, \epsilon_{ijc}^f) = \begin{cases} \begin{cases} \alpha_i \left(p\zeta - b - w_i\right) + \gamma_i + \epsilon_{ijc}^d & \text{drive} \\ 0 & \text{do not drive} \end{cases} & e_j, \phi_c = 1 \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

$\mathbf{p}$ is the price that ride share riders pay for rides. ζ is the portion of the payment that drivers keep after the app's commission. b represents expenses such as maintenance and foregone full-time employee benefits. w_i is the individual's outside wage. $\mathbf{p}\zeta - b - w_i$ therefore captures the dollar benefit of driving for ride share before financing costs. Observe that lower-income consumers obtain a greater economic benefit from ride share driving, other things equal. α_i scales this term from dollars to utility. γ_i captures the non-monetary net benefits of driving, such as the utility from having flexible hours. ϵ_{ijc}^d is an idiosyncratic shock that captures consumers with

[13] To support this assumption, I show in Appendix Section 6.5 that loan performance is unaffected by ride share entry, which suggests that lenders do not need to adjust interest rates in response to the driver's ride share decision.

observably similar economic benefits differing in their driving utility. For example, a parent may dislike driving at peak hours due to child care obligations.

Optimal consumer choice: The consumer's total indirect utility is the sum of her utility from own consumption, financing, and ride share driving. The consumer chooses which vehicle to own, how to finance it, and whether to drive for ride share. I assume that the consumer chooses financing and ride share driving after the purchasing decision, so that ϵ_{ijc}^{f} and ϵ_{ijc}^{rs} realize after the vehicle is acquired.[14] I characterize the solution to the financing and driving discrete choice problems before solving the overall vehicle choice problem.

Financing choice: Recall from (7) that a constrained consumer *must* obtain financing, and that an unconstrained consumer has the option to. The financing decision, with u_f^* denoting the optimized value function is therefore:

$$u_f^*(\theta_i, x_j', \epsilon_{ijc}^{f}) = -\alpha_i p_j + \begin{cases} -f(\theta_i, p_j, \epsilon_{ijc}^{f}) & l_i < p_j \\ \max_{f,\sim f}\{-f(\theta_i, p_j, \epsilon_{ijc}^{f}), 0\} & l_i \geq p_j \end{cases} \tag{9}$$

With the assumption that ϵ_{ijc}^{f} follows a type-one extreme value distribution, the probability that she obtains financing, and the expected indirect utility from the financing decision are can be analytically calculated and denoted as $p_f(\theta_i, x_j')$ and $u_f(\theta_i, x_j')$, respectively.

Ride share choice: Recall from (8) that a consumer with an eligible vehicle in a ride-share market has the option to be a ride share driver. The ride share decision, with u_{rs}^* denoting the optimized value function is:

$$u_{rs}^*(\mathbf{p}, \theta_i, x_j', \phi_c, \epsilon_{ijc}^{rs}) = \begin{cases} \max_{d,\sim d}\left\{\alpha_i\left(\mathbf{p}\zeta - b - w_i\right) + \gamma_i + \epsilon_{ijc}^{d}, 0\right\} & e_j, \phi_c = 1 \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

[14]The timing assumption of the ϵ shocks affords analytical tractability and are not critical for the main mechanism underlying the model, which is the relationship between the value of ride share driving and the need for finance related to driver demographics, which is independent of ϵ. The assumptions have basis in reality, where typically financing terms are negotiated after the consumer chooses a car, and where potential ride-share drivers are unable to evaluate the the idiosyncratic, non-monetary aspects of ride share driving, such as a preference for socialization with passengers, until they actually try driving.

With the assumption that ϵ_{ijc}^{rs} follows a type-one extreme value distribution, the probability that she drives for ride share[15] and her expected utility from the ride share can be analytically calculated and denoted as $p_{rs}(\mathbf{p}, \theta_i, x_j', \phi_c)$.

Vehicle choice: The consumer's indirect utility from choosing vehicle j is:

$$u(\mathbf{p}, \theta_i, x_j', \phi_c, \xi_{jc}, \epsilon_{ijc}^c) = u_c(\theta_i, x_j', \phi_c, \xi_{jc}) + u_f(\theta_i, x_j') + u_{rs}(\mathbf{p}, \theta_i, x_j', \phi_c) + \epsilon_{ijc}^c \tag{11}$$

Where u_f and u_{rs} are the expected utilities of the financing and ride share driving decisions. Her vehicle choice solves the following discrete choice problem:

$$\max_{j \in \{0,1,\cdots,J\}} \{u(\mathbf{p}, \theta_i, x_j', \phi_c, \xi_{jc}, \epsilon_{ijc}^c)\} \tag{12}$$

Again, assuming that ϵ_{ijc}^c follows a type-one extreme value distribution the ex-ante probability of consumer i choosing car j can be analytically calculated, which I denote as $p_{ijc} \equiv p(\mathbf{p}, \theta_i, x_j', \phi_c, \xi_{jc}; \{(x_k', \xi_{kc})\})$.

4.1.2 Aggregation

The preceding analysis characterized the decision problem for an individual consumer endowed with characteristics θ_i. To aggregate to market shares, I assume that consumer characteristics θ_i have the following distribution:

$$\theta_i = \bar{\theta} + (D_i - \bar{D})'\Pi + \Sigma\epsilon_i^\theta \tag{13}$$

$\bar{\theta}$ is an $n \times 1$ vector of preference means, constant across all consumers. D_i is a $d \times 1$ vector of individual demographics with mean $\bar{D}$, such as outside wage or whether she takes public transportation. Π is a $n \times d$ matrix mapping consumer demographics into consumer characteristics. Π governs, for example, how a consumer's outside wage impacts her availability liquidity and financing costs. Σ is an $n \times n$ variance-covariance matrix, and ϵ_i^θ is an $n \times 1$ iid vector of shocks, assumed to have a standard normal distribution. Individual demographics are themselves drawn from a market-level distribution, $D_i \sim F_c(D_i)$, measured directly in the data. This gives the key set of structural parameters to be estimated as $\Theta = (\bar{\theta}, \Pi, \Sigma)$.

[15]This probability can also be interpreted as the intensive margin of ride share driving.

Let $A_j(\mathbf{p}, x'_\cdot, \phi_c, \xi_{\cdot c}; \Theta)$ denote the set of consumer characteristics such that borrowers with these characteristics in market c choose car j. Then, the market share of vehicle j in market c given parameters Θ is

$$s_{jc}(\mathbf{p}, x'_\cdot, \phi_c, \xi_{\cdot c}; \Theta) = \int_{A_j} p(\mathbf{p}, \theta_i, x'_j, \phi_c, \xi_{jc}; \{(x'_k, \xi_{kc})\}) dF(\theta_i; \Theta) \tag{14}$$

Following the vehicle choice j, the conditional share of consumers taking up ride share and vehicle financing is:

$$s^{rs}_{jc}(\mathbf{p}, x'_\cdot, \phi_c, \xi_{\cdot c}|j; \Theta) = \int_{A_j} p^{rs}(\mathbf{p}, \theta_i, x'_j, \phi_c) dF(\theta_i|A_j; \Theta) \tag{15}$$

$$s^{f}_{jc}(x'_\cdot, \phi_c, \xi_{\cdot c}|j; \Theta) = \int_{A_j} p^{f}(\theta_i, x'_j) dF(\theta_i|A_j; \Theta) \tag{16}$$

4.1.3 Demand for ride share services and equilibrium

To close the model, I assume that aggregate demand for ride share services is:

$$\log q = \delta_0 - \delta_1 \log \mathbf{p} \tag{17}$$

Where q is the quantity of ride share services demanded, δ_0 is a constant, and δ_1 is the price elasticity of ride share services. I use estimates from Cohen et al. (2016) to calibrate δ_0 and the demand elasticity δ_1. Equilibrium $\mathbf{p}$ in the ride share service market equates demand and supply, $s^{rs}_{jc}(\mathbf{p}, x'_\cdot, \phi_c, \xi_{\cdot c}|j; \Theta) \times S$, where S is the number of households. Appendix Section 6.9 studies the sensitivity of my results to this this demand specification by varying price elasticity and the ride share commission rate.

Welfare: From the discrete choice framework described above, I calculate an individual driver's welfare directly as the ex-ante expected utility of (11); dividing by the individual's price sensitivity α_i expresses this quantity in dollar-equivalent terms. I calculate rider welfare from the demand equation (17). The following counterfactuals quantify welfare from riders and drivers but omit potential externalities such as increased congestion or pollution.

4.2 Estimation

Estimation strategy and identification: The estimation broadly follows the procedure in Berry et al. (1995) with several additions that utilize additional model predictions and micro-data on utilization, ride share driving, and financing. I instrument for vehicle price in the style of Hausman (1996), by using the Manufacturer Suggested Retail Price (MSRP), under the assumption that MSRP reflects cost shifters. Hence, the identifying variation comes from within-market differences in car prices rather than across market variation that could be contaminated by local demand shocks. For hourly ride share price p in the demand estimation, the ride share commission rate ζ, and the costs of ride share driving b, I take from Mishel (2018).[16] For the demand model, I use $\delta_1 = 0.57$ from Cohen et al. (2016).

The additional key parameters underlying the counterfactual results are (1) the average non-monetary value of driving, $\bar{\gamma}$, (2) the average cost of finance, $\bar{f}^0$, (3) the average liquidity availability, $\bar{l}$, and (4,5), how the cost of finance and availability of liquidity vary with consumer income, elements of Π. The average non-monetary value of driving, $\bar{\gamma}$, governs the total number of drivers that choose to enter ride share after accounting for the levels implied by the monetary benefits. Thus, conditioning on hte monetary benefits, modulating $\bar{\gamma}$ brings the model predicted driving share in line with the share observed in the data.

The key parameters governing a consumer's cost of finance, f_i^0 are the average cost of finance across all consumers, $\bar{f}^0$, and how f_i^0 varies with consumer income, which is an element in Π. Nearly all consumers in a high-income zip code possess sufficient liquidity to purchase an inexpensive vehicle without financing, but in the data, a non-zero share of consumers finance such cars. $\bar{f}^0$ governs this share: increasing $\bar{f}^0$ lowers the financing share while lowering $\bar{f}^0$ raises it. $\bar{f}^0$ is identified by matching the model predicted financing shares to those shares observed in the market. Additionally, the observed financing share of cheap vehicles varies across markets, and by adjusting the covariance between wage and financing costs in Π, the model can reproduce this relationship.

The key parameters governing a consumer's available liquidity, l_i, are the aver-

[16]This study itself combines the findings of several academic studies. As a baseline value, I take $p = \$22.06$ per hour, $\zeta = 0.75$, and $b = \$7.36$ per hour.

age liquidity available to all consumers, $\bar{l}$, and how l_i varies with consumer income, which is an element in Π. While financing *costs* are identified on the intensive margin of financing use for lower-priced vehicles, liquidity governs the extensive margin at which financing becomes necessary to purchase a vehicle. That is, a high value of $\bar{l}$ predicts that the price point at which financing shares approach 100% is high, while a low value of $\bar{l}$ predicts the opposite. By matching the overall level of this point, as well as its variation across markets, the estimation pins down average liquidity and how it varies with income.

Estimated parameters: Table 4 shows the model parameters, with Panel A showing the estimated parameters with their bootstrapped standard errors, and Panel B showing parameters from the literature. By dividing the parameters by the estimated price sensitivity, I translate them from utility values into dollar values. First, the mean log price sensitivity is roughly -2.9 with a standard deviation of 0.427. This corresponds to an average price sensitivity of 0.06, or a mean price elasticity of roughly 1.2, in line with estimates from Berry et al. (1995). Consistent with the reduced form evidence showing that upon ride share entry, high income households decrease utilization but sales remain constant, the estimation finds that customers place essentially no weight in the use value of a car, with the utility to car ownership coming from miles driven being essentially zero.

Financing costs are significant, estimated to be roughly $2 thousand dollars per year, roughly in line with typical financing terms on a car loan. These costs increase for low income households: A 1% decrease in outside option income corresponds roughly to a 0.45% increase in costs. The average household is estimated to have roughly $11 thousand of cash liquidity available to purchase a car, with higher income households having a significantly more. The estimated model closely matches cross-sectional and time-series patterns in financing rates. This fit is evaluated in more detail in Appendix Section 6.8.

4.3 Counterfactuals

I first use the estimated model to study how the gig economy's organization of production compares to that of a traditional taxi firm. I then study how financing costs affect the size, prices, distribution of capital, and employment in the gig economy, before considering technological advances that allow for an on-demand rental that makes possible the separation of ownership and driving in the gig economy.

4.3.1 Gig economy versus traditional production

I compare two mechanisms of capital allocation: First, a *traditional* allocation mechanism—a collection of competitive taxi companies. Here, firms obtain cars using finance, hire labor at the average outside option wage rate of current gig economy drivers adjusted for benefits and other non-monetary compensation[17]—and sell rides in a competitive market. I assume free entry of constant returns to scale firms that have financing costs equal to one quarter of the average financing cost for private individuals. This assumption reflects corporate borrowers having cheaper access to financial markets through expertise or by being able to spread idiosyncratic risk among a fleet of vehicles.[18] These assumptions pin down the equilibrium ride price as the price that drives profits to zero. I compare this to the *gig economy*, whose structure was described in detail above. Both scenarios fix the app's technology and regulatory environment, so that the only difference between these two structures is the capital allocation mechanism.

Figure 4 quantifies the differences between these mechanisms. Panel (a) shows ride share quantities, while Panel (b) shows ride prices. The black bars are the values for the traditional taxi firm structure, while the medium gray bars are the values for the gig economy. As compared to a traditional firm, there are roughly 11% more ride share rides in the gig economy, which are offered at a 26% discount. This reflects a shift outwards in the supply of rides, with prices and quantities determined by movement along the demand curve.[19] Panel (c) shows yearly net income from full-time driving, defined as the hourly driver income net of Uber commissions,

[17]For example, the cost of unemployment insurance that independent contractors must pay themselves. See Mishel (2018).

[18]Assuming that traditional firms face similar financing costs to private individuals, of course, will tend to magnify the differences that I calculate between these production structures, making gig economy production appear even more beneficial.

[19]Appendix 17 examines the effects for different demand elasticities. Qualitatively the results are very similar, with differing demand elasticities determining how the shift in supply is reflected in changes to quantities and prices.

driving expenses, and financing costs, assuming that drivers work eight hours per day for 250 days of the year. The figure shows that in the gig economy scenario, the net income of a driver is roughly \$13,000 per year, which is below the national minimum wage of \$7.25 per hour. Observe that by the assumption that taxi firms pay gig economy drivers their prevailing outside option wage, driver net incomes are similar, and differ only by the cost of financing, which gig economy drivers must pay themselves. This payment amount comes to roughly \$15,000 per year. Panels (e) and (f) show welfare effects of this organization of production compared to traditional firms. Panel (e), which isolates driver welfare, shows an increase in total driver welfare of roughly \$1 billion, while Panel (f), which includes both driver and rider welfare, shows an increase of more than \$2.5 billion.

The intuition is that a traditional firm acquires the car *solely* for the purposes of ride share driving, and so its cost must be justified through its income as productive capital alone. In contrast, in the gig economy or car share scenarios, car owners value cars both as productive capital and for their personal consumption. This means that they would be willing to own cars and supply them to the gig economy even when gig economy revenue alone does not justify acquiring the car. In consequence, there is a greater quantity of gig economy capital supplied. The car share mechanism goes one step further because it allows households with low financing costs to own cars and rent them to drivers rather than requiring that drivers with high financing costs own the cars themselves.

4.3.2 Financing costs in the gig economy

The reduced form results highlighted that its dependent on consumer finance was a potential cost of gig economy production, and conversely, easy access to credit may have played a role in its growth. In 2012, the year in which Uber opened its platform to independent drivers, the real interest rate on a 5-year auto loan was 2.1%, approximately three times lower than the historical average between 1980 and 2010 of 6.3%. To what extent did these low interest rates contribute to the gig economy's growth, and had rates been higher, would the gig economy be smaller and restricted to a different set of drivers?

To answer these questions, I allow $\bar{f}^0$, the average cost of finance, to vary between 0 and 10 times its baseline value. At three times its current level, this corresponds to the historical cost of finance; 10 times its current value corresponds to a case where effectively unavailable. Mirrored in the reduced form results, the model estimates imply that low-income individuals need credit to become ride share drivers. When credit is expensive, the after-finance net-income from driving for ride share decreases. In response, these finance-dependent individuals exit the market. Their exit decreases ride supply, reducing equilibrium ride quantity and increasing equilibrium ride price. This increases pre-finance driver income, and draws in higher-income drivers who can acquire cars without relying on finance.

Figure 5 quantifies this intuition and shows the results. Panel (a) shows ride share quantities, and panel (b) shows hourly ride prices. These figures show that as the costs of finance increase, ride quantity decreases and ride price increases. The relevant results for the gig economy are shown in dotted lines. At historical levels of finance costs, roughly three times current levels, the model estimates there would be roughly 7.5% fewer rides, and that prices would rise from roughly \$22 per hour to \$26 per hour. When financing prices become prohibitively high, quantities decrease by roughly 12.5% and prices rise to nearly \$30 per hour.

Changes in the price of finance lead to significant changes in driver income and demographics. The solid line in Figure 5 Panel (c) shows the yearly net income from full-time ride share driving. As financing costs increase, the net income from ride share driving initially decreases because higher financing costs directly increase the total costs of obtaining a car. As finance costs increase further, the net income of driving for ride share rises sharply, increasing to roughly 100% higher when finance is unobtainable. This occurs because as financing becomes more expensive, many low-income drivers are unable to obtain cars and therefore exit the market. As these drivers exit, the ride supply shifts inwards and ride prices increase. Increasing ride prices attract higher-income, less finance-dependent drivers to the market. The solid line in Panel (d), which shows the fraction of drivers who obtain their cars by borrowing, illustrates these selection dynamics. In the baseline scenario, nearly all drivers use financing, while as financing costs rise, the fraction of drivers who obtain financing drops to zero.

Panels (e) and (f) examine welfare. Initially, higher financing costs depress individual driver welfare, leading to a drop from roughly \$8 thousand per year to just about \$5 thousand per year. As financing costs increase, however, the dynamics that alter the driver composition lead to higher welfare for individual drivers as the drivers exposed to financing costs exit the market, and are replaced by non-finance reliant drivers who avoid the costs and receive higher wages. In terms of aggregate welfare, however, Panel (f) shows that higher financing costs unambiguously reduce welfare, both because drivers exit on the extensive margin, and consumers face lower quantities of rides at higher prices. With moderately more expensive finance, total welfare drops by more than \$4.5 billion relative to the baseline scenario.

Finally, the model quantifies the level of consumer finance costs that would make the gig economy structure hit parity with the traditional taxi firm structure in terms of quantities and prices. The intersection of the dotted line (gig economy) and solid line (traditional taxi) in Figure 5 Panels (a) and (b) address this question. By construction, higher consumer financing costs do not impact the supply of taxis in a traditional structure, but they do reduce the supply of ride share services in the gig economy. Prices and quantities intersect at slightly about three times the baseline—roughly at historical levels. Above this point, traditional taxis outperform the gig economy. This comparison illustrates the importance of the time period in which the gig economy arose, and also suggests an important comparative static across countries. In regions where the consumer finance system is less developed, one would expect traditional structures to be more common. In fact, this empirical regularity is largely true.

4.4 Car sharing technology

I finally consider the introduction of a hypothetical *car share* capital allocation mechanism and compare it to the *gig economy*. The gig economy requirement that drivers own their own cars induced a reliance on consumer finance and made quantities and prices were sensitive to financial conditions. In the *car share* capital allocation mechanism, I allow ride share drivers to rent idle cars from other households on demand.[20] As above, I assume that drivers earn their outside option under

[20]Importantly, the reduced form evidence and model estimates imply that using the car for ride share services does not reduce the vehicle's utility to its owner because consumers' own mileage

the gig economy scenario, with the residual income flowing to the vehicle owners.

Figure 4 Panel (a) shows that the introduction of a frictionless car share would increase quantities by roughly further 20% over the gig economy scenario, or by roughly 30% over the traditional taxi company. Panel (b) shows that a frictionless rental market would decrease prices from roughly the $22 per hour in the gig economy scenario to roughly $13 per hour. Panel (c) shows that these increases are not driven by offering low wages to drives. Panel (d) shows the financing rate of car owners whose vehicles are being used in the gig economy. Under the gig economy structure, this faction is nearly 100%, because owners and low-income drivers are one-in-the-same. Under the car share scenario, however, the fraction is closer to the unconditional financing rate of 40%. Panel (e) shows that the welfare of car-share car owners (the equivalent of drivers in the gig economy mechanism) increases by roughly $3.5 relative to the baseline and by $2.5 relative to the gig economy, and Panel (d) shows that the combined welfare of riders and ride share car owners increased by nearly $10 billion, or roughly $7.5 billion relative to the gig economy. Given that no such market has been widely adopted, an alternate interpretation of these large potential gains are as a lower bound for the size of the friction preventing its adoption.

Importantly, the introduction of a car share greatly reduces the reliance on finance. Figure 5 Panels (a) and (b) show that outcomes become insensitive to consumer finance costs despite car share, like the gig economy, using consumer durable goods for production. Panel (c) shows that driver net income rises with finance costs to bring higher income individuals into the gig economy, while in the case of car share, drivers are unaffected. Panel (d) shows the fraction of car owners using financing to sell vehicle services. In the case of car share, this fraction rapidly decreases to close to zero as financing costs rise and ownership shifts to less finance-reliant households. This shift is more rapid than in the gig economy because there is no countervailing force of higher wages demanding higher ride share compensation among the less finance-dependent houses. Panel (f) shows essentially no change in total welfare in the car share.

---

driven provides almost no direct utility above the convenience value of owning the car. In consequence, any utility loss from use will be small.

5 Discussion and Conclusion

The gig economy allows consumers to use durable consumption goods for production, and the ability to do so is a quantitatively meaningful benefit of the gig economy. This benefits come at the cost of making access to consumer finance a critical element in the gig economy production and participation. Access is especially critical for the low-income households who would like become ride share drivers but lack liquidity necessary to obtain a car.

When finance is available, auto sales, utilization, and employment increase following ride share entry. Conversely, in the absence of finance, these real effects are smaller. These reduced form facts motivated a structural model that quantified the net benefits of dual asset use to be roughly a \$3 billion increase in welfare for an industry with roughly \$30 billion in annual revenues as compared to a traditional taxi firm. However, without access to cheap consumer finance, these benefits disappear as finance-dependent low-income drivers are replaced with higher-income drivers who demand higher wages. The introduction of a frictionless rental market would increase consumer welfare by an additional \$7.5 billion annually and eliminate the gig economy's critical reliance on consumer finance by allowing vehicles to be owned by those who most value ownership and driven by those with the lowest opportunity cost of time.

These results highlight a novel feature of the gig economy that enables the flexible transition of consumer durables to capital goods that applies more broadly than in ride share. This feature helps explain, for example, why Uber, Lyft, and Airbnb—which all share this pattern of repurposing consumer durables for production—have successfully competed with traditional incumbents while companies like WeWork have not. This structure, however, relies on access to consumer finance, and my model suggests that their particular success was in large part aided by the low cost and easy access financial environment in which they arose.

Figure 1: Vehicle sales and auto loans after ride share entry by zip code income

Figure 1 shows the event study for auto sales, Panels (a), (c), and (e), and auto loans, Panels (b) and (d), and (f). Data are nationwide at the zip-quarter level from RL Polk (sales) and Equifax (loans). The left-hand side variable is log auto sales or log auto loans. Panels (a) and (b) show the results for all zip codes. Panels (c)-(e) split the sample by relative income of the zip code in the CBSA as of 2010, before ride share entry. Panels (c) and (d) show the effect for zip codes below the CBSA median income and Panels (e) and (f) show the effect for zip codes above the CBSA median income. Gray bars represent 95% confidence intervals, calculated by bootstrapping across CBSAs.

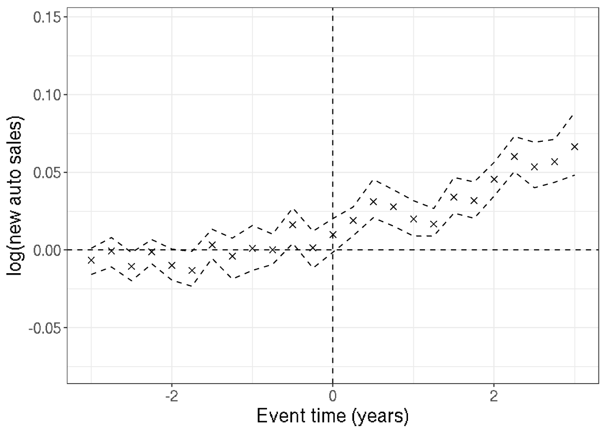

(a) New auto sales

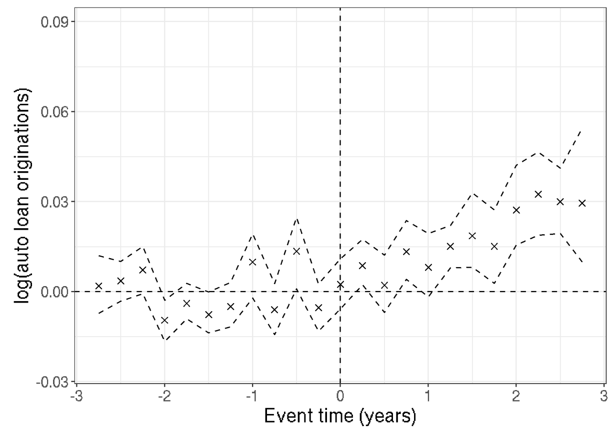

(b) New auto loan originations

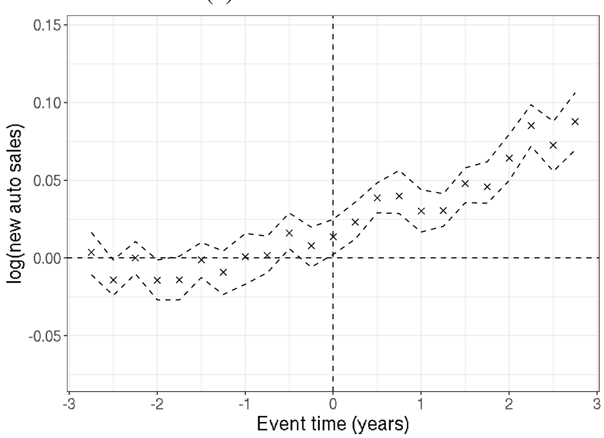

(c) Auto sales growth in low-income zips

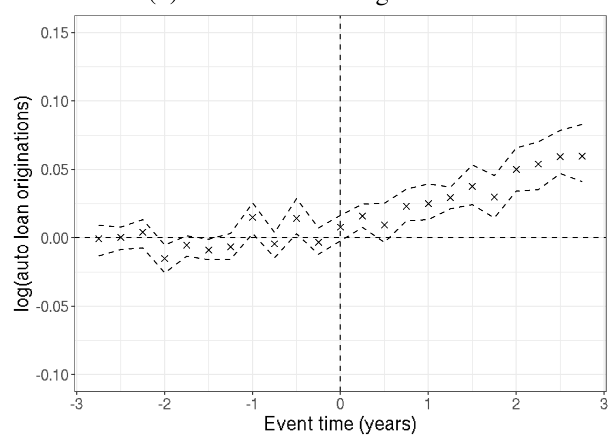

(d) Auto loan growth in low-income zips

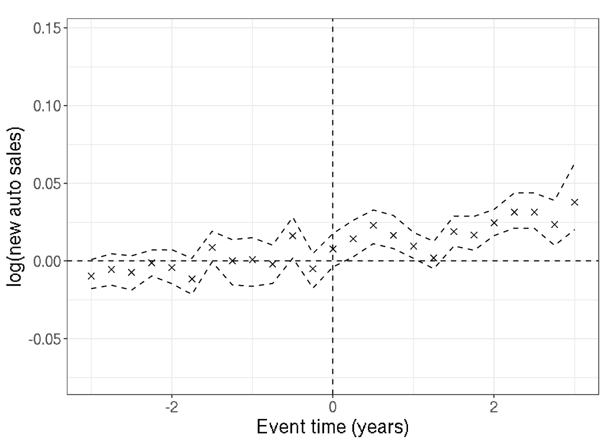

(e) Auto sales growth in high-income zips

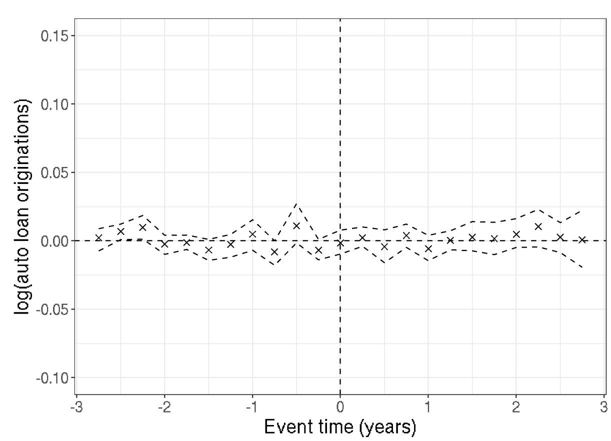

(f) Auto loan growth in high-income zips

Figure 2: Tax filings following ride share entry

Figure 2 shows the event study for tax filings, which uses nationwide zip-year level tax filing data from the IRS between 2010 and 2016. The left-hand side variable is log total tax filings within the zip code. Panel (a) shows the results using all tax filings. Panel (b) shows the results for tax filings with adjusted gross incomes equal to or below $25,000. Panel (c) shows the results for tax filings with adjust gross incomes above $25,000. Gray bars represent 95% confidence intervals.

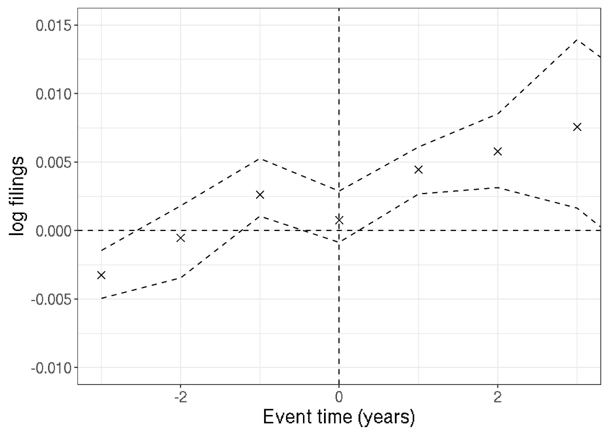

(a) Tax filings

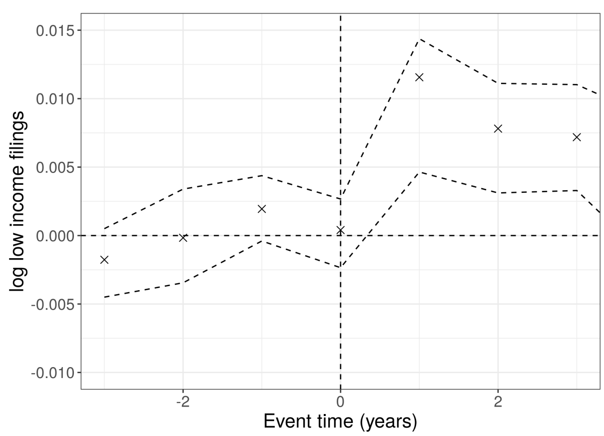

(b) Tax filings with AGI $\leq 25,000$

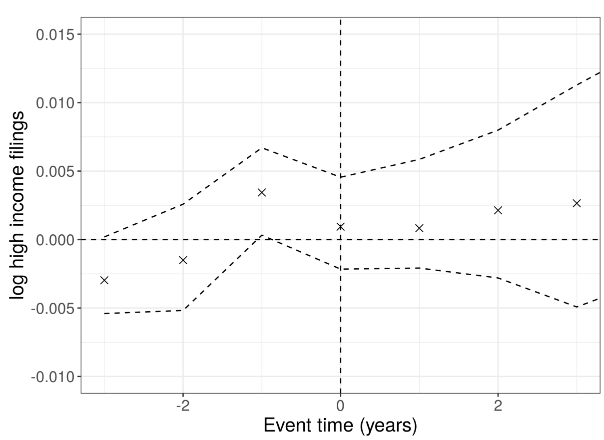

(c) Tax filings with AGI $> 25,000$

Figure 3: Bankruptcy disclosure and ride share financing

Figure 3 shows the impact of bankruptcy disclosure on auto lending and ride share entry. Panel (a) shows the quarterly raw probability of auto loan origination versus years since Chapter 7 bankruptcy filing. Panel (b) shoes the result of Regression (3) with separate coefficients for quarters after filing for bankruptcy, and includes zip-quarter fixed effects. Shaded regions are 95% confidence intervals. The darker regions represent the samples used in the ride share event study. Panels (c) and (d) show the cumulative probability of auto financing following ride share entry, with Panel (c) showing the results for borrowers with the bankruptcy flag during entry and Panel (d) showing the results for borrowers without the bankruptcy flag during entry.

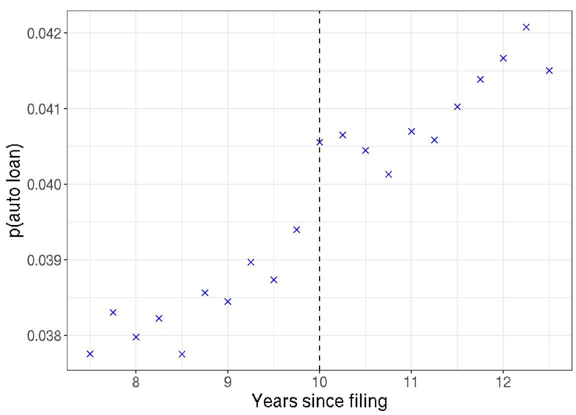

(a) Auto lending probability

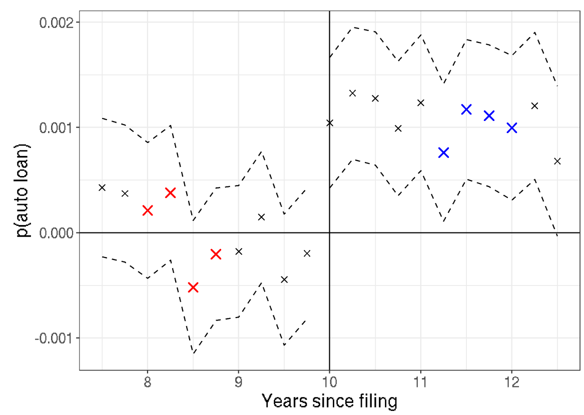

(b) Residualized auto lending probability

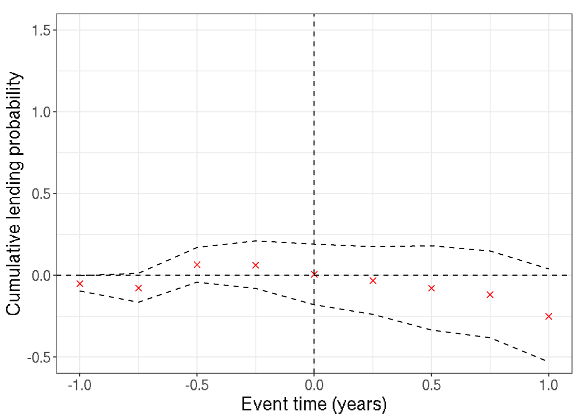

(c) Cumulative auto lending for constrained borrowers

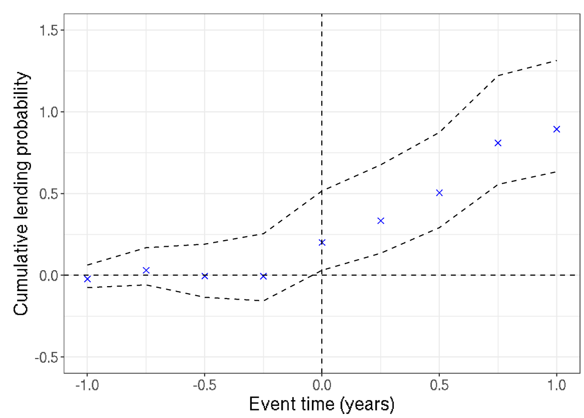

(d) Cumulative auto lending for unconstrained borrowers

Figure 4: Counterfactual systems of capital allocation

Figure 4 compares systems of capital allocation. The traditional firm scenario envisions firms owning financed cars and hiring drivers at minimum wage. The gig economy scenario is the current structure, where drivers work as independent contractors and own their cars. The car share scenario envisions individuals owning cars and renting them frictionlessly to drivers earning minimum wage. Panel (a) shows ride quantities relative to the baseline. Panel (b) shows ride prices per hour. Panel (c) shows driver net income. Panel (d) shows the percentage of drivers using finance. Panel (e) shows driver welfare relative to traditional firms, and Panel (f) shows driver and rider welfare relative to traditional firms.

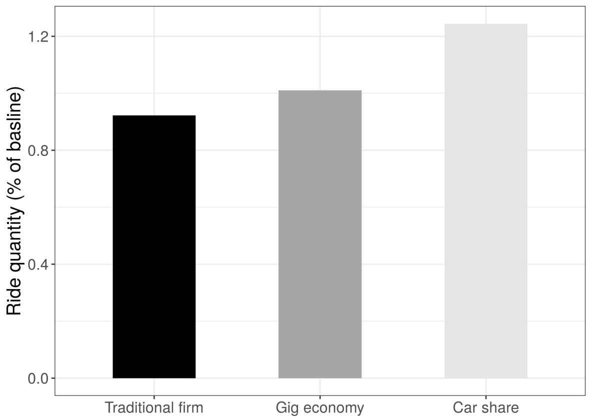

(a) Ride share ride quantities

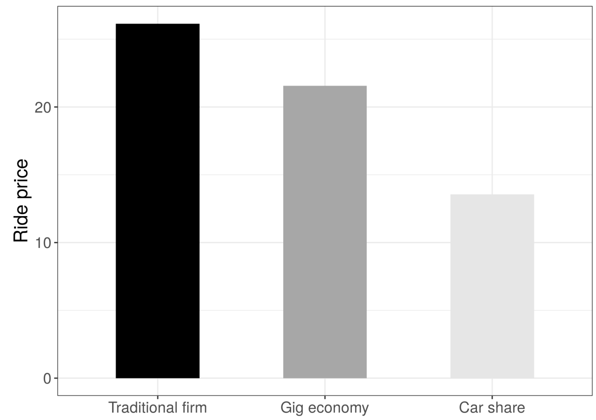

(b) Ride share ride prices

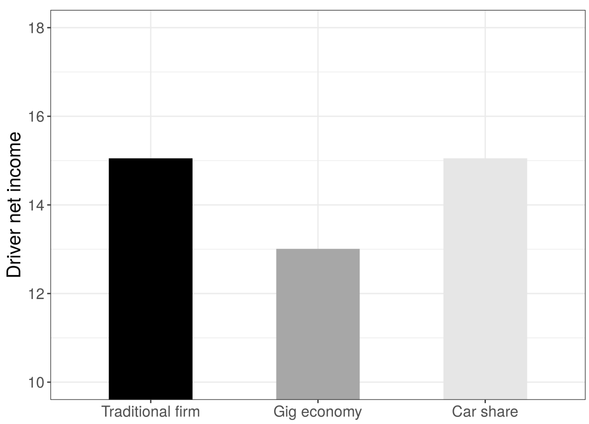

(c) Driver net income

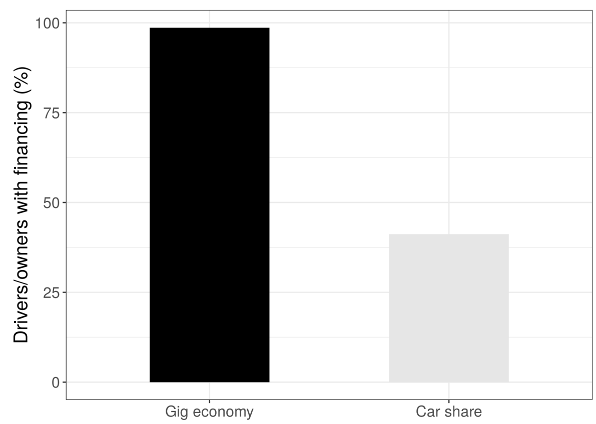

(d) Percent of drivers using finance

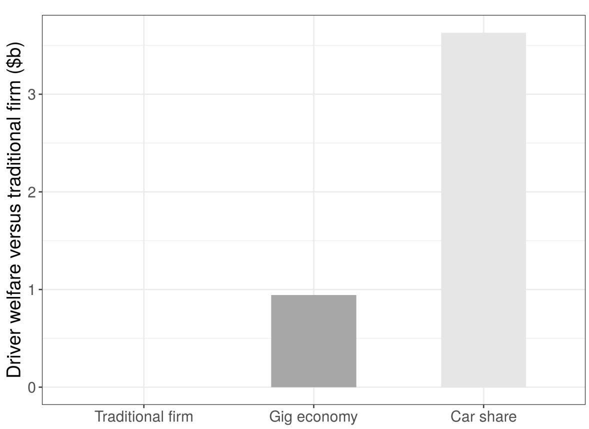

(e) Driver welfare

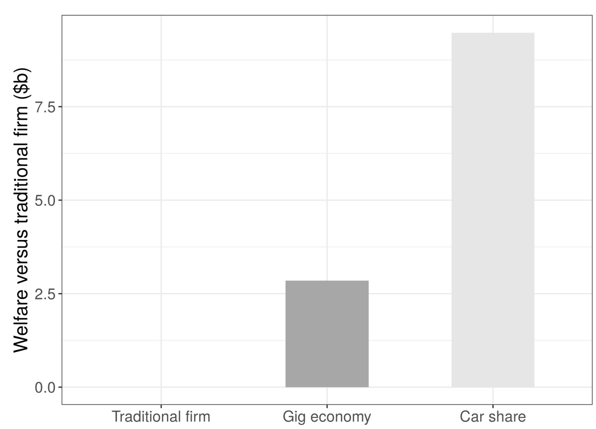

(f) Driver and rider welfare

Figure 5: Finance costs and counterfactual systems of capital allocation

Figure 5 shows how finance impacts ride share outcomes under different systems of capital allocation. A traditional firm owns financed cars and hires labor. The gig economy scenario is the current structure where drivers work as independent contractors and own cars. The car share scenario envisions individuals owning cars and renting them frictionlessly to drivers earning minimum wage. Panel (a) shows ride quantities relative to the baseline. Panel (b) shows ride prices per hour. Panel (c) shows driver net income. Panel (d) shows the percentage of drivers using finance. Panel (e) shows individual driver welfare in the gig economy. Panel (f) shows total welfare relative to the baseline scenario.

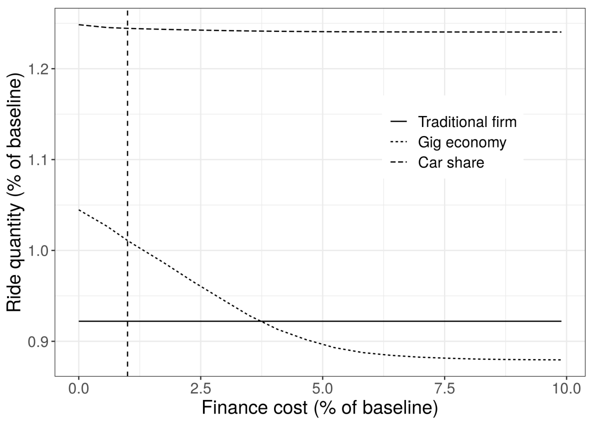

(a) Ride share ride quantities

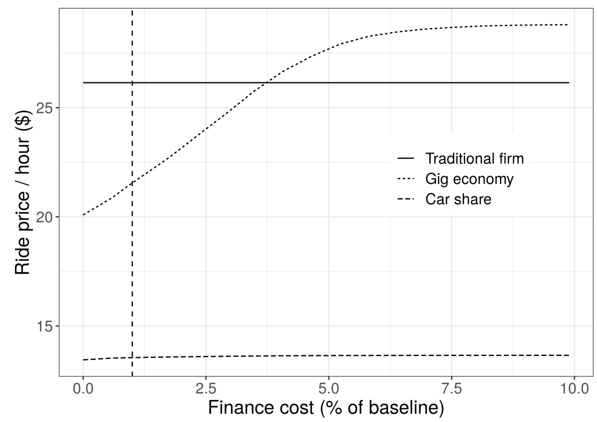

(b) Ride share ride prices

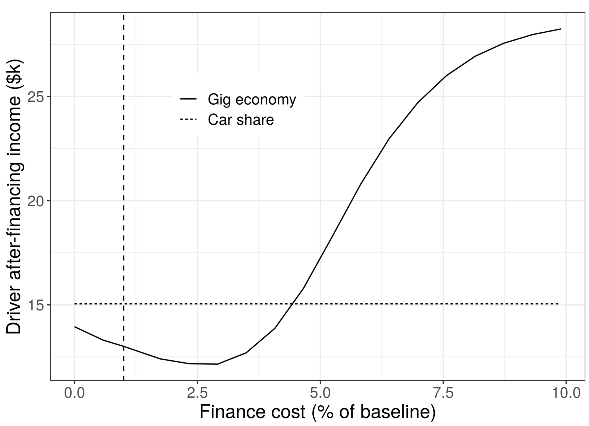

(c) Driver net income

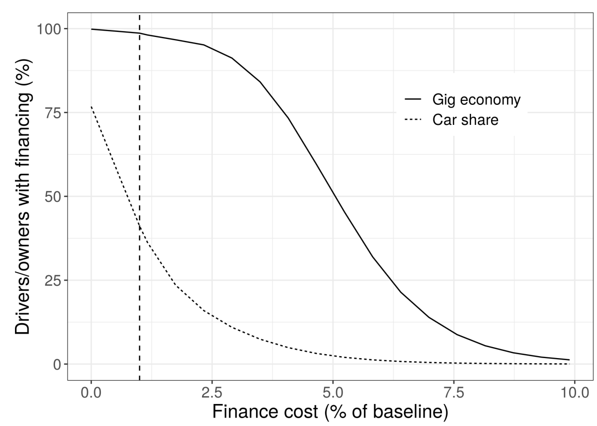

(d) Percent of drivers using finance

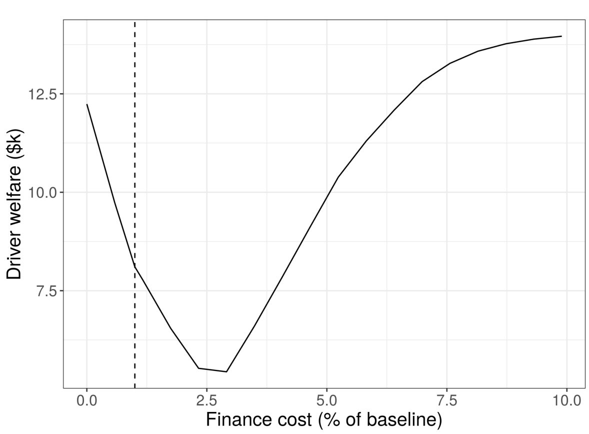

(e) Individual driver welfare

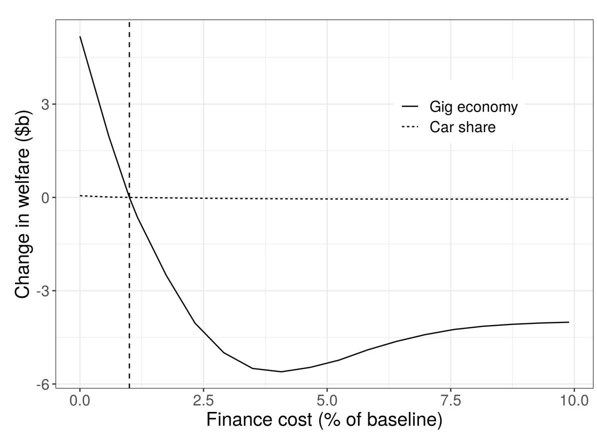

(f) Change in total welfare

Table 1: Nationwide auto sales and loan originations after ride share entry

Table 1 shows the results of Regressions (1) and (2) at the zip-quarter level:

$$\log Auto_{zt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Low\ Income_z + \gamma_z + \gamma_{Income,t} + \epsilon_{zt}$$

Panel A shows log new auto sales using nationwide RL Polk data between 2010 and 2016. Panel B shows log new auto loan originations using nationwide Equifax data between 2010 and 2016. All columns contain zip fixed effects. Column (1) is the standard difference-in-difference specification, where $Post$ signifies that Uber or Lyft has entered the zip's CBSA. Column (1) contains quarter fixed effects. Columns (2) and (3) interact $Post$ with whether the zip code is in the bottom 50% of median income within the CBSA, *Low income*, and whether the zip code is in the top 50% of share of transportation workers in the CBSA, *High transport share*. Columns (2) and (3) include low income and high transport share cross quarter fixed effects. Standard errors are clustered at the CBSA-quarter level.

Panel A: New auto sales, nationwide, RL Polk data

	Dependent variable:		
	Log sales		
	(1)	(2)	(3)
Post	0.016***	0.006***	0.004
	(0.002)	(0.002)	(0.003)
Post × Low income	-	0.020***	-
	-	(0.004)	-
Post × High transport share	-	-	0.025***
	-	-	(0.004)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr × Low income FE	N	Y	N
Qtr × High transport FE	N	N	Y
Observations	244,153	244,153	244,153
R^2	0.971	0.972	0.972

Panel B: New auto loan originations, nationwide, Equifax data

	Dependent variable:		
	Log new originations		
	(1)	(2)	(3)
Post	0.010***	−0.001	−0.002
	(0.002)	(0.002)	(0.002)
Post × Low wage	-	0.021***	-
	-	(0.003)	-
Post × High transport share	-	-	0.024***
	-	-	(0.003)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr × Low Wage FE	N	Y	N
Qtr × High transport FE	N	N	Y
Observations	244,153	244,153	244,153
R^2	0.979	0.979	0.979

Table 2: Employment and utilization after ride share entry

Table 2 shows the impact of ride share entry on employment and vehicle utilization. Panel A shows the result of regressing employment on ride share entry using IRS tax filing data between 2010 and 2016:

$$\log Filings_{bzt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times I(AGI_b < 25k) + FE + \epsilon_{bzt}$$

The regression is at the zip-year-tax bucket level. The left-hand side variable is the log number of tax filings. *Post* is an indicator for ride share entry. *AGI≤25k* is an indicator for whether the AGI bucket is below $25,000. All columns include tax bucket-year and tax bucket-zip fixed effects. Column (3) contains zip-year fixed effects. Standard errors are clustered at the CBSA-year level. Panel B shows the result of regressing vehicle utilization on ride share entry using DMV registration data from South Carolina between 2010 and 2016 at the vehicle-ownership spell-zip level:

$$MilesPerYear_{vzt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Eligible_v + \beta_3 + \beta_4 Post_{zt} \times Eligible_v \times Low\ income_z + FE + \epsilon_{vzt}$$

$MilesPerYear_{vzt}$ is the number of miles driven per year in thousands during the ownership spell. *Low income* is an indicator for whether the zip code is in the bottom 50% of median income in its CBSA. *Eligible* is an indicator for whether the vehicle is ride share eligible. All columns include zip $\times$ $Eligible$ and quarter $\times$ $Eligible$ fixed effects. Columns (4) and (6) include zip $\times$ quarter fixed effects. Standard errors are clustered at the CBSA-quarter level.

Panel A: Employment

	Dependent variable:		
	log filings		
	(1)	(2)	(3)
Post	0.006***	0.001	-
	(0.002)	(0.003)	-
Post × (AGI≤25k)	-	0.011**	0.011**
	-	(0.004)	(0.004)
(AGI<25k)×Year FE	Y	Y	Y
(AGI<25k)×Zip FE	Y	Y	Y
Zip×Year FE	N	N	Y
Observations	172,127	172,127	172,127
R^2	0.996	0.996	0.998

Panel B: Utilization

	Dependent variable:					
	Miles per year (k)					
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.175	−0.156	−1.033**	-	0.299	-
	(0.160)	(0.282)	(0.444)	-	(0.659)	-
Post × Low income	-	0.381	-	-	−1.518**	-
	-	(0.279)	-	-	(0.605)	-
Post × Eligible	-	-	1.350***	1.186**	−0.561	−0.922
	-	-	(0.469)	(0.543)	(0.714)	(0.905)
Post × Low income × Eligible	-	-	-	-	2.185***	2.372***
	-	-	-	-	(0.663)	(0.866)
Zip × Eligible FE	Y	Y	Y	Y	Y	Y
Qtr × Eligible FE	Y	Y	Y	Y	Y	Y
Zip × Qtr FE	N	N	N	Y	N	Y
Observations	129,215	129,190	129,215	129,215	129,190	129,190
R^2	0.036	0.036	0.036	0.158	0.036	0.158

Table 3: Bankruptcy disclosure and ride share financing

Table 3 shows the impact of bankruptcy disclosure on auto lending and ride share entry. Panel (A) shows the result of regressing auto loan origination on bankruptcy removal:

$$Origination_{izt} = \beta I(YearsSinceFiling \geq 10) + \gamma_{zt} + \epsilon_{izt}$$

$\geq$ 10 years is an indicator for whether the bankrupt consumer's bankruptcy was 10 or greater years in the past. Columns (1)–(4) vary the observation window around the ten-year removal window, expanding the window from $\pm$0.25, 1.00, 1.50, and 2.50, respectively. *$\geq$10 years* is an indicator for whether the filing occurred more than 10 years from the observation date. Panel (B) shows the result of regressing origination on ride share entry interacted with bankruptcy removal:

$$Origination_{izt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Constrained_i + \gamma_{gt} + \gamma_{gz} + \gamma_{zt} + \epsilon_{izt}$$

Post is an indicator for ride share entry. Constrained is an indicator for whether the consumer is in the constrained (has bankruptcy filing on credit report) group. Columns (1)–(3) use a bankruptcy flag removal window of one year; (4)–(6) use a window of 0.50 years. Columns (1) and (4) include only the treatment indicator; Columns (2)–(3) and (5)–(6) include the interaction with flag removal. Columns (1)–(2) and (4)–(5) include zip-flag group and date-flag group fixed effects; Columns (3) and (6) additionally include zip-quarter fixed effects. Standard errors, in parentheses, are clustered at the CBSA-flag group level.

Panel A: Bankruptcy flag and auto originations

	Dependent variable:				
	P(auto loan) (%)				
	(1)	(2)	(3)	(4)	(5)
Window (years)	$\pm$0.25	$\pm$0.50	$\pm$1.00	$\pm$1.50	$\pm$2.50
$\geq$ 10 years	0.132***	0.149***	0.132***	0.136***	0.115***
	(0.032)	(0.021)	(0.015)	(0.012)	(0.010)
Zip-Time FE	Y	Y	Y	Y	Y
Observations	2,052,307	4,021,994	7,799,010	11,303,095	17,332,333
R^2	0.228	0.146	0.091	0.068	0.049

Panel B: Bankruptcy flag removal and ride share entry

	Dependent variable:					
	P(auto loan) (%)					
	Window = $\pm$ 1 year			Window = $\pm$0.50 years		
	(1)	(2)	(3)	(4)	(5)	(6)
Post_{zt}	0.085	0.243***	-	0.093	0.255**	-
	(0.059)	(0.083)	-	(0.081)	(0.109)	-
$\text{Post}_{zt} \times$ Constrained	-	-0.316^{**}	-0.317^{**}	-	-0.310^{*}	-0.326^{*}
	-	(0.128)	(0.142)	-	(0.171)	(0.197)
Zip-Group FE	Y	Y	Y	Y	Y	Y
Date-Group FE	Y	Y	Y	Y	Y	Y
Zip-Time FE	N	N	Y	N	N	Y
Observations	1,920,408	1,920,408	1,920,408	1,073,389	1,073,389	1,073,389
R^2	0.019	0.019	0.073	0.028	0.028	0.115

Table 4: Model parameters

Table 4 shows the model parameters. Panel B shows the parameters estimated from consumer demand BLP estimation. $\bar{\beta}$ is the mean parameter value. Π_{wage} is how the parameter varies with demographics. Σ is the standard deviation of the parameter shock. Bootstrapped standard errors are sown in parentheses. Panel B shows parameters taken from the literature and their sources.

Panel A: Estimated parameters

Parameter	Description	$\bar{\beta}$ (SE)	Π_{wage} (SE)	Σ (SE)
Convenience and utilization				
β_0^{pre}	Intercept (pre)	-3.841 (0.025)	0.373 (0.133)	0.486 (0.146)
β_0^{post}	Intercept (post)	-3.295 (0.038)	0.819 (0.124)	0.486 (0.146)
m^{pre}	Utilization (k miles) pre	12.335 (0.640)	-	0.493 (1.058)
m^{post}	Utilization (k miles) post	11.560 (0.328)	-	0.493 (1.058)
$\log \alpha$	Log price sensitivity	-2.962 (0.257)	-	0.427 (0.134)
Car characteristics				
β_m	Use utility	0.000 (0.658)	-	-
β_e	Eligibility value	1.638 (0.027)	-	-
β_{hp}	HP value	-0.238 (0.029)	-	-
β_{mpg}	MPG value	-0.034 (0.028)	-	-
Financing and driving				
f_0^{pre}	Financing cost (pre)	0.683 (0.137)	-0.445 (0.081)	0.503 (0.101)
f_0^{post}	Financing cost (post)	0.859 (0.172)	-0.445 (0.081)	0.503 (0.101)
l^{pre}	Liquidity (pre)	10.961 (0.080)	1.854 (0.075)	0.316 (0.050)
l^{post}	Liquidity (post)	10.723 (0.647)	1.854 (0.075)	0.316 (0.050)
γ	Drive utility	-0.181 (0.208)	-	0.262 (0.111)
T1EV Variances				
$\phi_{finance}$	Financing draw variance	0.331 (0.134)	-	-
ϕ_{drive}	Driving draw variance	0.003 (0.008)	-	-

Panel B: Other parameters

Parameter	Description	Value	Source
p	Ride price	\$22.06 per hour	Mishel (2018)
b	Hourly non-finance costs	\$7.36 per hour	Mishel (2018)
ζ	Uber commission	0.75	Mishel (2018)
ξ	Booking fee	\$1.55	Mishel (2018)
δ_1	Demand elasticity	0.57	Cohen et al. (2016)

References

Barrios, J. M., Hochberg, Y. V., and Hanyi Yi, L. (2018). The cost of convenience: Ridesharing and traffic fatalities. *Working Paper*.

Benetton, M. (2018). Leverage regulation and market structure: An empirical model of the uk mortgage market. *Available at SSRN 3247956*.

Benetton, M., Fantino, D., et al. (2018). Competition and the pass-through of unconventional monetary policy: evidence from tltros. Technical report, Bank of Italy, Economic Research and International Relations Area.

Benjaafar, S., Ding, J.-Y., Kong, G., and Taylor, T. (2018). Labor welfare in on-demand service platforms.

Benmelech, E., Meisenzahl, R. R., and Ramcharan, R. (2017). The real effects of liquidity during the financial crisis: Evidence from automobiles. *The Quarterly Journal of Economics*, 132(1):317–365.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890.

Buchak, G., Matvos, G., Piskorski, T., and Seru, A. (2018a). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics*.

Buchak, G., Matvos, G., Piskorski, T., and Seru, A. (2018b). The limits of shadow banks.

Buera, F. J., Kaboski, J. P., and Shin, Y. (2011). Finance and development: A tale of two sectors. *American Economic Review*, 101(5):1964–2002.

Campbell, J. Y. (2006). Household finance. *The journal of finance*, 61(4):1553–1604.

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.

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

Cook, C., Diamond, R., Hall, J., List, J. A., Oyer, P., et al. (2018). The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *Working Paper*.

Cook, C., Diamond, R., and Oyer, P. (2019). Older workers and the gig economy. In *AEA Papers and Proceedings*, volume 109, pages 372–76.

Cramer, J. and Krueger, A. B. (2016). Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–82.

Di Maggio, M., Egan, M., and Franzoni, F. A. (2019). The value of intermediation in the stock market. *Available at SSRN 3430405*.

Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A., and Yao, V. (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review*, 107(11):3550–88.

Dobbie, W., Goldsmith-Pinkham, P., Mahoney, N., and Song, J. (2016). Bad credit, no problem? credit and labor market consequences of bad credit reports.

Donaldson, J. R., Piacentino, G., and Thakor, A. V. (2018). Household debt and unemployment.

Egan, M., Hortaçsu, A., and Matvos, G. (2017). Deposit competition and financial fragility: Evidence from the us banking sector. *American Economic Review*, 107(1):169–216.

Fraiberger, S. P. and Sundararajan, A. (2017). Peer-to-peer rental markets in the sharing economy.

Greenwood, R. and Scharfstein, D. (2013). The growth of finance. *Journal of Economic Perspectives*, 27(2):3–28.

Grunewald, A., Lanning, J., Low, D., and Salz, T. (2019). Auto dealer loan intermediation: Consumer

behavior and competitive effects. Technical report, mimeo.
Hall, J. V., Horton, J. J., and Knoepfle, D. T. (2017). Labor market equilibration: Evidence from uber. *URL http://john-joseph-horton. com/papers/uber_price. pdf, working paper*.
Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In *The economics of new goods*, pages 207–248. University of Chicago Press.
Herkenhoff, K., Phillips, G., and Cohen-Cole, E. (2016). The impact of consumer credit access on employment, earnings and entrepreneurship.
Horton, J. J. and Zeckhauser, R. J. (2016). Owning, using and renting: Some simple economics of the" sharing economy".
Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448.
Hurst, E. and Stafford, F. P. (2004). Home is where the equity is: Mortgage refinancing and household consumption. *Journal of money, Credit, and Banking*, 36(6):985–1014.
Jayaratne, J. and Strahan, P. E. (1996). The finance-growth nexus: Evidence from bank branch deregulation. *The Quarterly Journal of Economics*, 111(3):639–670.
Kaplan, S. N. and Zingales, L. (1997). Do investment-cash flow sensitivities provide useful measures of financing constraints? *The quarterly journal of economics*, 112(1):169–215.
Lenzu, S. and Manaresi, F. (2017). Do marginal products differ from user costs? micro-level evidence from italian firms.
Melzer, B. T. (2011). The real costs of credit access: Evidence from the payday lending market. *The Quarterly Journal of Economics*, 126(1):517–555.
Mian, A. and Sufi, A. (2011). House prices, home equity-based borrowing, and the us household leverage crisis. *American Economic Review*, 101(5):2132–56.
Mian, A., Sufi, A., and Verner, E. (2017). How do credit supply shocks affect the real economy? evidence from the united states in the 1980s.
Midrigan, V. and Xu, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *American economic review*, 104(2):422–58.
Mishel, L. (2018). Uber and the labor market.
Musto, D. K. (2004). What happens when information leaves a market? evidence from post-bankruptcy consumers. *The Journal of Business*, 77(4):725–748.
Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.
Ostrovsky, M. and Schwarz, M. (2018). Carpooling and the economics of self-driving cars.
Philippon, T. (2015). Has the us finance industry become less efficient? on the theory and measurement of financial intermediation. *American Economic Review*, 105(4):1408–38.
Piskorski, T. and Seru, A. (2018). Mortgage market design: Lessons from the great recession. *Brookings Papers on Economic Activity*, 2018(1):429–513.
Razeghian, M. and Weber, T. A. (2016). To share or not to share: Adjustment dynamics in sharing markets.
Zinman, J. (2010). Restricting consumer credit access: Household survey evidence on effects around the oregon rate cap. *Journal of banking & finance*, 34(3):546–556.

6 Appendix for Online Publication

This appendix contains supplemental material. Section 6.1 details the data sources used in the paper, presents summary statistics, and compares DMV registrations to new auto sales. Section 6.2 shows regressions concerning the intensive and extensive margins of ride share entry. Section 6.3 studies the ex-ante distribution of vehicle ownership in the economy.

Section 6.4 contains an extended analysis of the impact of ride share ownership on new vehicle registrations and vehicle liens, using highly detailed DMV registrations data. This analysis complements the main findings of the paper to show that ride share eligible vehicles in low-income zip codes saw registrations and financing increase following ride share entry. Section 6.5 contains the analysis of auto loan performance following ride share entry. Section 6.6 studies the impact on tax filings, broken out by business and wage or salary income. Section 6.7 provides information on the experimental design of the bankruptcy flag analysis and considers regional level measures of financial constraints.

Section 6.8 examines the fit of the estimated model. Section 6.9 examines the sensitivity of the counterfactuals to different demand and platform parameters. Section 6.10 contains placebo analyses of the main empirical results of the paper.

6.1 Data

My paper brings together a number of datasets described here. Summary statistics are shown in Table A1.

Uber and Lyft Data: I use two sources of data on Uber and Lyft entry, one public and one propriety from Uber. The public dataset lists the dates and locations of ride share entry for Uber and Lyft. I hand collect this data from news reports and press releases. Uber and Lyft entry occurs at the city level, which I map to core based statistical areas (CBSAs). Entry begins in 2010 and continues in a staggered manner through 2017.

The proprietary dataset provided by Uber is at the Uber market-month level. An Uber market corresponds approximately to a CBSA. The dataset provides the number of active drivers in each market and at each month. A driver is active if he or she picks up at least one passenger in given market during the month.

RL Polk auto sales data: I use new vehicle sales data provided by RL Polk. This data is collected through dealerships and state car registration databases. The dataset provides the number of new vehicle sales at the zip-month level between 2010 and 2016. It does not include used vehicle sales.

Department of Motor Vehicles (DMV) Data: I use vehicle registration data from the South Carolina, Washington, and Indiana DMVs. Car owners must register their cars yearly with their state's DMV. A registration requires an inspection and a nominal fee. Therefore, this registration data provides a comprehensive measure of the active automobile stock in a state across time. My dataset spans 2010 to 2016 and includes the vehicle identification number (VIN), zip code, and month of registration. Additionally, South Carolina records odometer readings when the vehicle changes ownership. Washington records whether there is a lien against the vehicle, which signifies whether the vehicle was acquired through a secured loan. Appendix Table A2 compares the DMV new registrations datasets and the RL Polk new sales datasets to insure that they are consistent and finds that they are.

NHTSA and FuelEconomy.gov Databases: The VIN recorded in the DMV registrations dataset contains an identifier that merges with the National Highway Traffic Safety Administration's (NHTSA) database. This database provides detailed physical attributes of the car, including model year, manufacturer, engine displacement and horsepower, number of doors, and body type. I further merge this dataset with fuel economy and emissions data from FuelEconomy.gov. This merge provides a complete picture of the physical attributes of every registered vehicle in the DMV dataset.

Equifax loan originations data: I use data from Equifax to obtain the number of loan auto originations and auto loan performance at the zip-quarter level. The dataset separately reports originations from banks and from captive auto lenders. A captive auto lender is a non-bank lender connected to a dealership, like GM Financial. The data covers 2010 to 2016.

TransUnion borrower-level data: I use borrower-level data from TransUnion. This data is a 10% nationwide sample of all individuals in the United States on which TransUnion has information. It provides the individual's zip code, his or her loan originations across many asset classes, and information on his or her past bankruptcy filings. In my paper, I extract the auto borrowing behavior between 2009 and 2016 of all individuals in the sample who declared Chapter 7 bankruptcy before 2012.

IRS Summary of Income: The IRS provides zip-year level information on tax filings. This data reports the total number of individuals and households filing tax returns in a zip-year. It further breaks down the number of filers into adjusted (AGI) brackets. For example, it provides the number of filers in a zip-year with an AGI falling below $25,000.

Demographic Data: Finally, I use standard demographic variables from the 2010 United State Census and American Community Surveys between 2010 and 2016. I use individual data on car ownership rates and household income, as well as zip- and CBSA-level data on population, income, unemployment, public transportation use, mobile broadband access, worker commutes, education, race, age, and others.

6.2 The timing and location of ride share entry

Uber entered San Francisco in 2010, with Lyft following shortly thereafter. Both services subsequently expanded rapidly across the united states. Figure A3 Panel (a) shows the staggered entry of ride share across markets by month. Panel (b) shows the total number of Uber drivers in the United States between 2012 and 2016. Panel (c) plots the number of drivers per resident centered around the time of Uber entry. This number grows to roughly 0.15% two and a half years after entry, and roughly 0.35% four years after entry.

I next study the determinants of ride share entry along the extensive and intensive margins with the following cross-sectional regression run at the CBSA level:

$$\begin{aligned} Entered_c &= X_c'\beta + \epsilon_c \\ YearsToEntry_c &= X_c'\beta + \epsilon_c \end{aligned} \tag{18}$$

$Entered_c$ is a dummy variable taking the value one if Uber or Lyft entered the market between 2010 and 2017. This captures entry on the extensive margin. $YearsToEntry_c$ is a continuous variable equal to the number of years after 2010 that Uber or Lyft first entered the market. This captures entry on the intensive margin. X_c' is a vector of CBSA characteristics including log population, percentage of households with mobile broadband access, and, importantly for this paper, percentage of households with vehicles, among other controls detailed below. These variables are included in both levels and changes. The level variables are as of 2010;[21] the change variables are calculated as three year changes. Table A4 shows the results. Columns (1) and (2) study the extensive margin of entry; columns (3) and (4) study the timing of entry among entered markets. Columns (1) and (3) include only the main covariates of interest; Columns (2) and (4) include additional covariates.[22]

Across all specifications, there is a strong association between population and entry. A 1% larger population is associated with a 0.20% greater probability of platform entry, and among entered cities, a 1% larger population predicts entry roughly 0.7 years earlier. Additionally, mobile broadband access predicts 5.5 to 6 years earlier entry. Variables concerning vehicle ownership and financial access are not statistically significant predictors of entry. These results show that Uber and Lyft enter cities with large markets: large cities where consumers have the mobile broadband access necessary to use the ride share applications.

6.3 Ex-ante determinants of vehicle ownership

While the rate of vehicle ownership in the United States is high, ownership is heterogeneous across the income distribution. This is important for ride share because ride share driving pays a low wage, and if the individuals who might become ride share drivers do not have cars, they will need to obtain them. Using individual data from the 2010 ACS, I calculate the number of cars per adult household member and run the following household-level regression:

$$CarsPerHH_i = \sum_b \gamma_b I(Income_i \in Bin_b) + \gamma_c + \epsilon_i \tag{19}$$

$CarsPerHH_i$ is household i's vehicles per adult household member. $I(Income_i \in Bin_b)$ is an indicator for whether household i's income falls within income quantile Bin_b. γ_c is a CBSA fixed effect. These coefficients of interest, γ_b, show how car ownership varies non-parametrically across the income distribution relative to others in the CBSA. Figure A3 Panel (d) plots these coefficients. Vehicle ownership rates are correlated with household income. Households earning at or below the full-time Uber income, indicated by the vertical dashed line, have roughly 0.20 fewer cars per adult household member than households earning the median income, suggesting that prior to ride share entry, ride share capital is distributed away from the potential drivers.

To study a richer set of determinants of vehicle ownership and sales, I run the following zip-level regression:

$$\begin{aligned} Y_z &= X_z'\beta_1 + X_c'\beta_2 + \epsilon_z \\ Y_z &= X_z'\beta_1 + \gamma_c + \epsilon_z \end{aligned} \tag{20}$$

The left-hand side variable, Y_z, is either the average number of vehicles per household in the zip code, or new auto sales per hundred residents. Zip-code vehicle ownership comes from the 2010 census; new auto sales comes from 2010 RL Polk data. X_z' and X_c' are vectors of zip and CBSA controls, respectively. γ_c is a CBSA fixed effect. The primary covariates of interest are zip code income and zip code unemployment, although many others are included in the specification.[23] The specification

[21] Except for the case of mobile broadband access, which is first available as of 2013.

[22] Those additional covariates are: Log income, Δ Income, log house price, Δ house price, % households with mortgages, Δ households with mortgages, % households with high-speed Internet, and Δ households with high-speed Internet.

[23] These include: At the zip-code level, average commute time, log population, percent using public transit, percent working outside home, average age, percentage with a bachelors degree, percentage white, and percentage of of auto loans

with CBSA fixed effects, in particular, directly examines whether it is the high- or low-income zip codes within a city that own or buy cars.

Table A5 shows the results. Columns (1) and (2) use vehicles per household as the left-hand side variable; Columns (3) and (4) use new sales per hundred residents as the left-hand side variable. Columns (1) and (3) include several CBSA-level controls, while Columns (2) and (4) include CBSA fixed effects. Focusing first on Column (2), the results show that positive labor market conditions within the CBSA are strongly associated with vehicle ownership: households in zip codes with greater unemployment rates, both in absolute terms and relative to the CBSA have significantly fewer cars per household and have fewer auto sales per household. Similarly, households in a zip code with 1% greater income relative to the CBSA have 0.24 more cars per household and 0.24 auto purchases per hundred residents.

6.4 DMV Registrations and Liens

DMV registrations: The analysis in the body of the paper focused on new auto sales aggregated at the zip-code quarter level. It did not differentiate among the types of cars being bought. In this section, I exploit detailed DMV registration data from South Carolina, Indiana, and Washington. This dataset contains all vehicle registrations in the state. Registrations are required to be renewed annually. These data provide not only a measure of the active capital stock in an area at a given time, but also include a unique Vehicle Identification Number (VIN) for each car. A car's VIN, when merged with other publicly available data described earlier, provides a detailed description of its physical characteristics, such as its manufacturer, model age, and body type.

While this data is restricted in geographical scope, it offers several advantages that complement the preceding nationwide analysis. First, it measures registrations of *all* vehicles, and not just new vehicle sales. Second, it allows me to study how ride share differentially impacts vehicles of different types. This is particularly useful because ride share services place restrictions on which vehicles are eligible to be driven on their platforms: To be *eligible*, generally speaking, a vehicle must be no more than 15 years old, have four doors, and be a sedan, SUV, or minivan. To the extent that individuals are acquiring cars to drive for ride share, increases in auto sales or registrations should be concentrated among vehicles that are eligible to be driven on the platform. This data allows me to exploit within-CBSA and within-zip variation to test whether this is in fact true.

My outcome variable of interest is $\log Regs_{mezt}$, the log of the number of new registrations of manufacturer m of eligibility status e, in zip code z at time t. A unit of observation is, for example, the number of Uber or Lyft-eligible Nissans in zip code 60615 in January 2014. A registration is a new registration if the VIN was not present in the zip code in the previous period.[24]

I begin with an event studying testing how new registrations of eligible vehicles in low wage zip codes respond differentially to ride share entry:

$$\log Regs_{mezt} = \sum_{\tau=-4}^{4} \beta_\tau I(t - ET_z = \tau) \times Low\ income_z \times Eligible_{me} + \gamma_{tme} + \gamma_{zme} + \epsilon_{mezt} \tag{21}$$

$Eligible_{me}$ is a zero-one indicator for whether the vehicle is eligible. $Low\ income_z$ is a zero-one

and mortgages that are subprime. At the CBSA level: log income, percentage of households with high-speed Internet, percentage of households with mobile broadband, and log population.

[24]This measure omits vehicles changing hands within the zip code, but does measure the same owner moving to a new zip code and registering her car in the new zip code.

indicator for whether the zip code is in the bottom half of income among zips in the CBSA. The coefficient on the interaction of $I(t - ET_z = \tau) \times Low\ income_z \times Eligible_{me}$ shows the differential effect of ride share entry on new auto registrations for eligible vehicles in low-income zip codes. γ_{tme} is a time $\times$ manufacturer $\times$ eligibility fixed effect, which controls for how eligible and ineligible vehicle registrations of a particular manufacture change systematically nationwide. This absorbs, for example, the effect of Toyota introducing a particularly popular car during the sample period that happens to be ride-share eligible, but whose popularity is unrelated to ride share entry. γ_{zme} is a zip $\times$ manufacturer $\times$ eligibility fixed effect, which controls for certain makes and models being particularly popular in a given region. For notational convenience, I omit the separate interaction terms of event time $\times$ $Low\ income_z$ and event time $\times$ $Eligible_{me}$, but these terms are included in the regression as estimated.

Figure A6 Panels (a) and (b) show the coefficients on event time, together with 95% confidence intervals. Panel (a) shows the coefficient on event time $\times$ $Low\ Income_z \times Eligible_z$, picking up the differential effect of ride share entry on ride share-eligible vehicles in low-income zip codes. Panel (b) shows the baseline coefficient on event time $\times$ $Low\ Income_z$ without regards to the vehicle's eligibility. Neither panel exhibits pre-trends. Panel (a) shows that registrations of ride share eligible vehicles in low-income zip codes increase substantially following ride share entry, while Panel (b) shows that ride share entry has no effect on ineligible vehicles in low-income zip codes. These figures, taken together, show that increases in new auto registrations occur precisely among the populations and vehicle types consistent with their owners entering ride share driving.

To confirm these results, I run the following quadruple difference specification at the manufacturer-eligibility-zip-time level:

$$\begin{aligned}\log Regs_{mezt} = {} & \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Eligible_{me} + \beta_3 Post_{zt} \times Low\ income_z \\ & + \beta_4 Post_{zt} \times Eligible_{me} \times Low\ income_z + t \times Eligible_{me} \times Low\ income_z \\ & + \gamma_{mez} + \gamma_{met} + \gamma_{zt} + \epsilon_{mezt}\end{aligned} \tag{22}$$

Variables in the regression are as defined previously in equations (2) and (21). The specification mirrors the previous specifications with the addition of $t \times Eligible_{me} \times Low\ income_z$. These are a set of linear time trends varying by eligibility and income bucket. These differential time trends rule out, for example, the possibility that households in low-income zip codes increase their purchases of eligible vehicles at a faster rate than high-income zip codes for reasons unrelated to ride share entry. γ_{zt} is a zip $\times$ time fixed effect, not included in all specifications. This fixed effect rules out ride share entry being correlated with local economic effects that impact the purchase of eligible and ineligible vehicles simultaneously. For example, this rules out ride share entry itself leading to economic growth that induces more local car sales. Identification in Regression (22) comes from variation in vehicle eligibility $\times$ zip $\times$ time. That is, the staggered entry provides zip-time variation while vehicle eligibility provides within zip-time variation.

Table A7 Panel A shows the results. Column (1) shows no effect in overall new vehicle registrations. Column (2), which includes the $Post \times Low\ income$ interaction, finds a statistically significant effect of roughly 1% more new registrations in low-income zip codes following ride share entry. This is statistically indistinguishable from the differential result for low-income zip codes found in the nationwide study in Table 1 Panel A Column (2). Column (3) includes the $Post \times Eligible$ interaction, as well as the triple interaction of $Post \times Low\ income \times Eligible$. The regression shows roughly a 2% differential increase in new registrations of ride share eligible vehicles in low-income zip codes. Column (4) repeats the specification in Column (3) but includes zip $\times$ time fixed effects and finds the same differential result.

As before, I run a placebo test by randomizing the locations that receive ride share entry. This randomization allows me to check that these specifications are not picking up spurious trends related to vehicle eligibility and income that are unrelated to ride share entry. The results of this placebo test are shown in Table A21 Panel A and show no significance.

DMV liens data: I now exploit the detailed DMV registrations data from Washington state, which indicates whether a particular vehicle has a lien attached to it. The presence of a lien indicates that the vehicle is securing a loan, so this measures whether the vehicle was obtained with financing. I study whether ride-share eligible vehicles are more likely to be financed following ride share entry, and whether this is particularly true in low income zip codes. An affirmative answer would suggest that as the quantity of eligible vehicles increases in low-income zip codes, these new vehicle acquisitions are effected through increased borrowing.

To study this question, I run the following regression at the vehicle registration level:

$$\begin{aligned} HasLien_{vzt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Eligible_v + \beta_3 Post_{zt} \times Low\ income_z \\ + \beta_4 Post_{zt} \times Eligible_v \times Low\ income_z + t \times Eligible_v + FE + \epsilon_{vzt} \end{aligned} \tag{23}$$

$HasLien_{vzt}$ is a zero-one indicator variable for whether there is a lien attached to the vehicle at the time of registration. $t \times Eligible_v$ adds separate linear time trends for eligible and ineligible vehicles interacted with high and low income indicators, intended to capture varying trends in likelihood of obtaining financing that varies by vehicle eligibility and zip income. The term FE collects various fixed effects, the details of which are described below. The coefficient of interest is β_4, which captures the differential impact of ride share entry on eligible vehicles in low-income zips. Table A7 Panel B shows the results.

Column (1) shows the overall treatment effect and contains zip, $Low\ income \times$ month, and make-model fixed effects. It finds essentially no effect in the likelihood of a vehicle being financed. Column (2) adds the post times low-income interaction with the same fixed effects and again finds no effect. Column (3) adds the triple interaction with vehicle eligibility and shows that ride share-eligible vehicles in low-wage zip codes are significantly more likely to have liens—roughly 2% more likely—following ride share entry as compared to ride-share ineligible vehicles in low-income zip codes. Column (4) adds zip $\times$ times quarter fixed effects to absorb any changes in local economic conditions that may correlate with ride share entry. With these fixed effects, all variation arises from the within-zip code differential impact between eligible and ineligible vehicles. These results show an increase in the likelihood of obtaining a lien of roughly 1.2% following ride share entry for eligible vehicles relative to ineligible vehicles. A placebo test randomizing which cities receive treatment, shown in Appendix Table A21, Panel B, shows no effects.

6.5 Loan performance after ride share entry

I study the impact of ride share entry on loan performance. To do so, I repeat specification (1) and (2) where the outcome of interest is loan performance:

$$\begin{aligned} &Performance_{zt} = \beta Post_{zt} + \gamma_t + \gamma_z + \epsilon_{zt} \\ &Performance_{zt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Low\ Income_z + \gamma_z + \gamma_{Income,t} + \epsilon_{zt} \end{aligned}$$

Where $Performance_{zt}$ is the the percentage of loan originations that are sixty or more days delinquent within one year of origination. Observations at time t in zip z represent how loans originated at time t at zip z performed over the following year. Table A8 shows the results.

Column (1) shows the raw impact of ride share entry on performance. The coefficient of -0.0003 means that auto loans following ride share entry are 0.03 percent less likely to become seriously delinquent. Column (2) shows the interaction with low income zip codes, and finds that while delinquencies decrease overall, delinquencies in low-income zip codes decrease by a slightly smaller amount. The same holds true in Column (3) which looks at the differential effect of high transport share zip codes.

The conclusion of this analysis is that ride share entry is associated with very small, statistically significant improvements in loan performance, but the effects are slightly more muted in low income zip codes. Given the small effect, it is difficult to conclude much from these results aside from the fact that ride share entry, and the subsequent consumer borrowing, does not lead to large increases in loan default. In other words, the borrowing increases following ride share entry appear to be manageable debt increases among borrowers, at least given how I measure loan performance here.

6.6 Business versus salary and wage income

The IRS summary of income data reports those tax filings that include salary earnings and those that include business income. Because during the sample period, Uber and Lyft classified their drivers as independent contractors, for tax purposes, drivers' income is classified as business income rather than salaries or wages. As a robustness check, I test whether increases in tax filings following Uber's entry come from increases in filings with reported wages and salaries, or from increases in filings with reported business income. Unlike earlier results breaking out by wage buckets, I consider total filings across all wage buckets. The specification is as follows:

$$\log Filings_{bzt} = \beta_1 Post_{zt} + FE + \epsilon_{bzt} \tag{24}$$

The regression is at the zip-year level. Only zips that eventually receive treatment are in the regression. The left-hand side variable is the log number of tax filings. *Post* is an indicator variable taking the value 1 after ride share has entered. All columns include year and zip fixed effects.

Table A9 shows the results. Column (1) considers the total number of filings of any time. Column (2) considers the number of filings reporting salary and wage income. Column (3) considers the number of filings reporting business income. All columns include zip and year fixed effects. All columns include These results show that wage and salary filings increase by a small but statistically insignificant amount following ride share entry. These filings increase by 0.3%. In contrast, filings reporting business income increase by 1.3%, which is economically meaningful and statistically significant.

6.7 Experimental design and zip-level evidence of financial constraints

This section details the empirical design of the bankruptcy flag removal analysis and provides further evidence of financial constraints at the zip code level, and their interaction with ride share. Figure A10 details the timeline of the event study.

Zip-level financial constraint analysis: The zip-level empirical strategy mirrors regression (2):

$$Y_{zt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Constrained_z + \gamma_z + \gamma_{Income,t} + \epsilon_{zt}$$

Rather than splitting by zip-level income, I split by the ex-ante bank share and the ex-ante consumer finance delinquency rate, captured by $Constrained_z$. I begin by studying loan originations. Figure A11 shows the event studies for loans. Panels (a) and (b) show the sample splits by low- and high-auto loan bank share, respectively. This figure shows that auto lending in unconstrained zip codes with a low 2010 bank share has no pre-trends, with auto lending increasing rapidly once ride share enters. Constrained zip codes with a high 2010 bank share, on the other hand, see no change once ride share enters. Similarly, Panels (c) and (d) show the sample splits by low- and high-2010 consumer finance delinquency rates, respectively. As before, the unconstrained zip codes with low delinquency rates show flat pre-trends and then a sharp increase in auto lending following ride share entry, while the constrained zip codes with high delinquency rates show no effect from ride share entry.

Next, I test whether these effects on lending carry through to auto sales and employment. I use the same specifications, replacing the left-hand side variable with log auto sales and log tax filings. The results for auto sales are shown in Figures A12 for auto sales and A13 for tax filings. In both figures, Panels (a) and (b) measure financial constraints with the 2010 bank share of auto lending, and Panels (c) and (d) measure financial constraints with the 2010 consumer finance default share. Panels (a) and (c) are the unconstrained zip codes while panels (b) and (d) are the constrained zip codes. These figures show a significant increase in auto sales and tax filings in the unconstrained zip codes (Panels (a) and (c)) and no effect in the constrained zip codes (Panels (b) and (d)).

Regressions confirm the results. Because high banks share and high default share are correlated with zip income, I include $Low\ income \times Post$ as an additional control. Table A14 Panel A shows the results for auto loans. Column (1) uses bank share as the measure of financial constraints; Column (2) uses default share as the measure of financial constraints. In both cases, there is a significant increase in auto lending after ride share entry, but confirming the graphical results in Figure A11, in high bank share zip codes and high default share zip codes, the effect is zero or reversed. To summarize, in both cases, borrowing increases in the unconstrained zip codes but remains constant in the constrained zip codes.

Finally, Table A14 Panel B shows the results for auto sales and tax filings. Columns (1)–(2) show the results for auto sales, and Columns (3)–(4) show the results for employment. Columns (1) and (3) use bank share as the measure of financial constraints, while Columns (2) and (4) use consumer defaults as the measure of financial constraints. Each regression shows a significant baseline effect among the unconstrained zip codes, with log auto sales and employment increasing by roughly 2% and 1% respectively. However, in zip codes with financial constraints, these effects vanish. A placebo test, shown in Appendix Table A24, which randomizes the location of ride share treatment and shows no effect.

6.8 Evaluation of model fit

I evaluate the fit of this critical relationship between income and financing behavior in Figure A15. Panels (a)–(c) plot the actual and estimated relationship between wage and financing in three samples of the data. In each panel, the dashed line is the model prediction and the solid line is the data. Panel (a) shows the results over the entire sample; Panel (b) shows the results before ride share entry; Panel (c) shows the results post ride-share entry. Observe that the model closely captures both the overall level of financing as well as both the direction and magnitude of the negative relationship between financing and wage particularly in the post-ride share subsample. Panel (d) plots the percentage of individuals who actually obtain financing versus the model's prediction, and shows a close relationship between the two.

6.9 Model sensitivity to demand assumptions

I close the model with a simple demand for ride share rides specification. In the body of the paper I assume that the price elasticity of demand is 0.57 from Cohen et al. (2016), and following Mishel (2018), that ride share apps charge a fixed 25% commission with the remainder of the ride price going to the driver, i.e., $1 - \zeta = 0.25$. In this section, I perform a number of robustness checks around the demand elasticity and commission level. In particular, I test how sensitive my quantitative and qualitative findings are to these assumptions.

Varying demand elasticity: I vary the demand elasticity point estimate by a factor of $\pm 25\%$ around its current value. That is, I consider it taking the values $\{0.42, 0.57, 0.72\}$. A16 shows how the financing cost results depend on demand elasticity. Unsurprisingly, quantity responses are much larger in the case of elastic demand and much smaller in the case of inelastic demand. The other comparative statics are largely unaffected. Figure A17 shows how outcomes change among the three capital allocation structures for different elasticities. This figure shows no qualitative differences, with the only quantitative difference occurring in ride share quantities. Higher elasticities expand the differences between the capital allocation mechanisms in terms of quantities, while lower elasticities compress them.

Varying ride share commission: I vary the commission level by $\pm 10\%$. That is, I allow $1 - \zeta \in \{0.15, 0.25, 0.35\}$. Figure A18 shows how the financing cost results depend on the ride share commission. Not surprisingly, a lower commission increases the overall quantity and price levels across financing costs. Commissions do not, however, alter the relationship between financing costs and these outcomes, nor do they impact levels or sensitivities to financing costs of driver income or demographics. Figure A19 shows how the commission level impacts outcomes across the capital allocation structures. For the traditional firm, there is mechanically no effect because the firm simply keeps the rider's payment and pays the driver a fixed wage. In the other cases, there are small quantitative differences that do not alter any qualitative conclusions in comparing these capital allocation mechanisms.

6.10 Placebo analyses

This section performs placebo tests of the main specifications. The placebo treatment is constructed by retaining the actual dates of ride share entry, but randomly assigning them across MSAs. The placebos are the zip-level sales and loans analysis, Table A20; the vehicle-level stock and liens analysis, Table A21; the IRS employment placebo, Table A22; the vehicle utilization placebo, Table A23, and the credit supply placebo, Table A24. In all cases, the placebo tests return negative results, as would be expected if the timing and location of ride share entry drives the results in the analysis.

Table A1: Summary statistics

Table A1 shows select summary statistics for the main datasets used in the paper. Panel A shows data for the nationwide datasets, RL Polk, which shows vehicle sales, and Equifax, which shows auto loan originations. Panel B shows vehicle registrations data from the DMVs of South Carolina, Indiana, and Washington.

Panel A: RL Polk and Equifax data

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Sales	567,874	132.367	305.593	0	10	172	34,019
New originations	567,874	126.713	166.496	0	13	186	2,889
Outstanding loans	567,874	2,266.832	2,944.440	1	235	3,392	32,959
Sales per capita	567,874	0.011	0.017	0.000	0.006	0.013	0.988
New loans per capita	567,874	0.012	0.007	0.000	0.008	0.015	0.273
Outstanding loans per capita	567,874	0.208	0.091	0.0001	0.156	0.242	0.998

Panel B: DMV vehicle registrations data

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
New registrations	74,761	279.814	584.161	0	5	305	9,789
New eligible registrations	74,761	229.941	496.153	0	3	244	9,264
Percent eligible	63,972	0.789	0.140	0.000	0.747	0.852	1.000

Table A2: New registrations and new sales comparison

Table A2 compares new sales data from RL Polk and new registrations data from select states' DMVs between 2010 and 2016. Columns (1)-(2) regress total sales on new registrations. Columns (3)-(4) regress log sales on log new registrations. Columns (5)-(6) regress sales per capita and new registrations per capita. Odd columns have no fixed effects; even columns have quarter and zip fixed effects. Standard errors are shown in parentheses.

	Dependent variable:					
	Sales		Log sales		Sales / capita	
	(1)	(2)	(3)	(4)	(5)	(6)
New registrations	0.353***	0.339***	-	-	-	-
	(0.002)	(0.004)	-	-	-	-
Log new registrations	-	-	0.989***	0.261***	-	-
	-	-	(0.002)	(0.009)	-	-
New registrations per capita	-	-	-	-	0.400***	0.416***
	-	-	-	-	(0.004)	(0.004)
Constant	−7.050***	-	−1.214***	-	−0.002***	-
	(0.821)	-	(0.010)	-	(0.0001)	-
Zip FE	N	Y	N	Y	N	Y
Qtr FE	N	Y	N	Y	N	Y
Observations	29,626	29,626	29,626	29,626	29,626	29,626
R^2	0.623	0.868	0.890	0.970	0.293	0.755

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Figure A3: Timing of ride share entry and ex-ante vehicle ownership

Figure A3 shows Uber and Lyft entry through time. Panel (a) shows the number of new markets that Uber or Lyft enter by month, using data collected from press releases and news reports. Entry is defined by the first time that either Uber or Lyft begin operations in the market. Panel (b) shows the total number of Uber drivers (not including Lyft drivers) in the United States, using Uber data. Panel (c) shows the number of Uber drivers per resident in Uber markets around the time that Uber enters, using Uber data. Panel (d) shows the number of vehicles per household member over eighteen years of age versus household income after removing CBSA fixed effects. The vertical dashed line represents the average full-time equivalent ride share driving income of $23,500. See Mishel (2018). Data are from the 2010 ACS. Gray bars represent 95% confidence intervals.

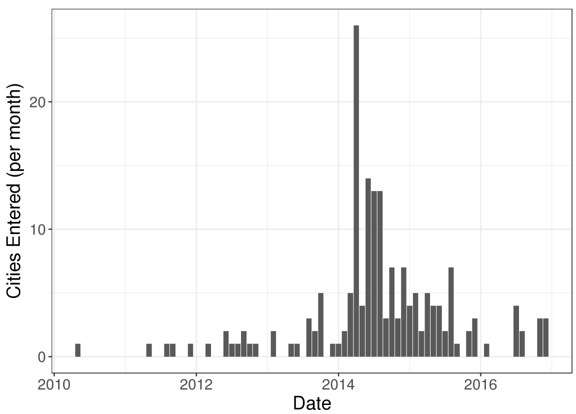

(a) New markets entered by month

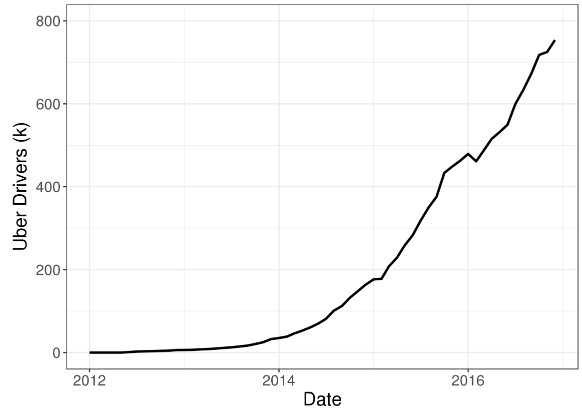

(b) Total Uber drivers in the United States

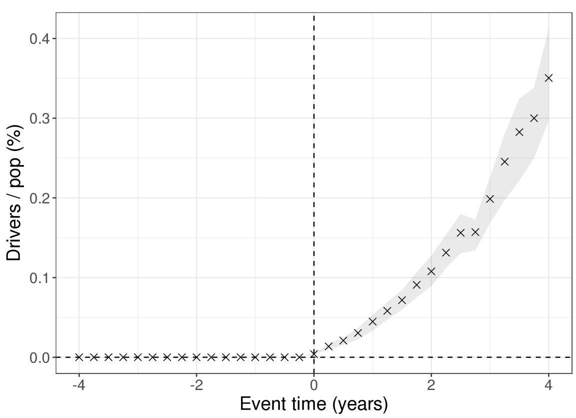

(c) Driver share of population around entry

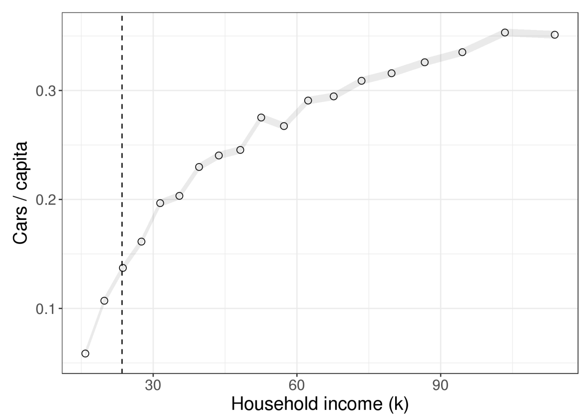

(d) Vehicles per household versus household income

Table A4: Timing of ride share entry

Table A4 shows the determinants of ride share entry as specified in Regression 18. All specifications are cross-sectional and on the CBSA level using data as of 2010 and 2013. Columns (1) and (2) study the extensive margin of entry. The left-hand side variable is a dummy taking the value 1 for CBSAs that ride share enters and 0 otherwise. Columns (3) and (4) study the intensive margin of ride share entry by restricting the sample to entered CBSAs. The left-hand side variable is the number of years that pass since 2010 before ride share enters. *Log population* is the log CBSA population. *HH with mobile broadband* is the percentage of households with mobile broadband access. *HH with vehicles* in the percentage of households with vehicles. *Bank share of auto financing* is the percentage of auto loans originated by banks (as opposed to captive lenders). Δ controls are changes between 2010 and 2013. Columns (2) and (4) include unreported additional controls: Log income, Δ Income, log house price, Δ house price, % households with mortgages, Δ households with mortgages, % households with high-speed Internet, and Δ households with high-speed Internet. Standard errors are in parentheses.

	Dependent variable:			
	Entered		Years to entry	
	(1)	(2)	(3)	(4)
log population	0.229^{***}	0.235^{***}	-0.799^{***}	-0.739^{***}
	(0.021)	(0.025)	(0.071)	(0.088)
Δ population	1.450	0.815	−4.209	−0.498
	(1.261)	(1.436)	(4.508)	(5.173)
HH with mobile broadband (%)	0.754^{**}	0.925^{**}	-5.905^{***}	-5.150^{***}
	(0.332)	(0.377)	(1.202)	(1.457)
Δ % HH with mobile broadband	0.064	0.159	-2.486^{*}	-2.861^{*}
	(0.319)	(0.341)	(1.412)	(1.500)
HH with vehicles (%)	0.364	1.266	−0.619	−2.092
	(1.000)	(1.073)	(3.285)	(3.580)
Δ % HH with vehicles	−1.322	−0.678	8.664	7.715
	(1.326)	(1.447)	(5.872)	(6.216)
Bank share of auto financing	0.009	0.098	−0.955	−0.204
	(0.186)	(0.218)	(0.691)	(0.861)
Other controls	N	Y	N	Y
Observations	470	460	215	214
R^2	0.320	0.334	0.523	0.542
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A5: Vehicle stocks and sales before ride share entry

Table A5 shows the result of regressing ownership and sales per household on zip- and CBSA characteristics:

$$Y_z = X_z'\beta_1 + X_c'\beta_2 + \epsilon_z Y_z = X_z'\beta_1 + \gamma_c + \epsilon_z$$

The regression is at the zip-code level as of 2010. Columns (1) and (2) use cars per household from the census as the left-hand side variable; Columns (3) and (4) use new auto sales per household from RL Polk as the left-hand side variable. Columns (1) and (3) include CBSA controls; Columns (2) and (4) include CBSA fixed effects.*Unemployment rate (zip)* is the 2010 unemployment rate at the zip-code level; *Log income (zip)* is the 2010 zip code median income; *Average commute time (zip)* is the average commute time at the zip level. *CBSA* variables are as defined in Table 4. All columns include the following unreported zip-level controls: Log population, % commuting on public transit, % working outside the home, average age, % with bachelors degrees, % white, % of auto loans and residential mortgages that are subprime. Columns (1) and (3) include log population as a control at the CBSA level. Standard errors are in parentheses.

	Dependent variable:			
	Cars per household		Sales per capita	
	(1)	(2)	(3)	(4)
Unemployment rate (zip)	-0.015^{***}	-0.016^{***}	-0.016^{***}	-0.014^{***}
	(0.001)	(0.001)	(0.002)	(0.003)
Log income (zip)	0.240^{***}	0.244^{***}	0.358^{***}	0.236^{***}
	(0.012)	(0.012)	(0.047)	(0.052)
Average commute time (zip)	0.014^{***}	0.014^{***}	-0.005^{***}	-0.006^{***}
	(0.0003)	(0.0003)	(0.001)	(0.002)
Log income (CBSA)	-0.003	-	-0.279^{***}	-
	(0.017)	-	(0.069)	-
HH with high-speed Internet (CBSA)	-0.472^{***}	-	0.377^{***}	-
	(0.025)	-	(0.104)	-
HH with mobile broadband (CBSA)	0.725^{***}	-	-1.106^{***}	-
	(0.039)	-	(0.159)	-
CBSA FE	N	Y	N	Y
Other Controls	Y	Y	Y	Y
Observations	16,935	16,935	16,982	16,982
R^2	0.722	0.790	0.110	0.145
Note:				*p<0.1; **p<0.05; ***p<0.01

Figure A6: New vehicle registrations following ride share entry

Figure A6 shows the event study for vehicle registrations, which uses DMV data available for Washington, South Carolina, and Indiana. The left-hand side variable is log auto registrations in the zip code. Panel (a) shows the difference-in-difference coefficient for ride-share eligible cars in low income zip codes. Panel (b) shows the difference in difference coefficient for ride-share ineligible cars in low income zip codes. A ride share eligible car is no more than 15 years old, has four doors, and is a sedan, SUV, or minivan. Gray bars represent 95% confidence intervals, calculated by bootstrapping across CBSAs.

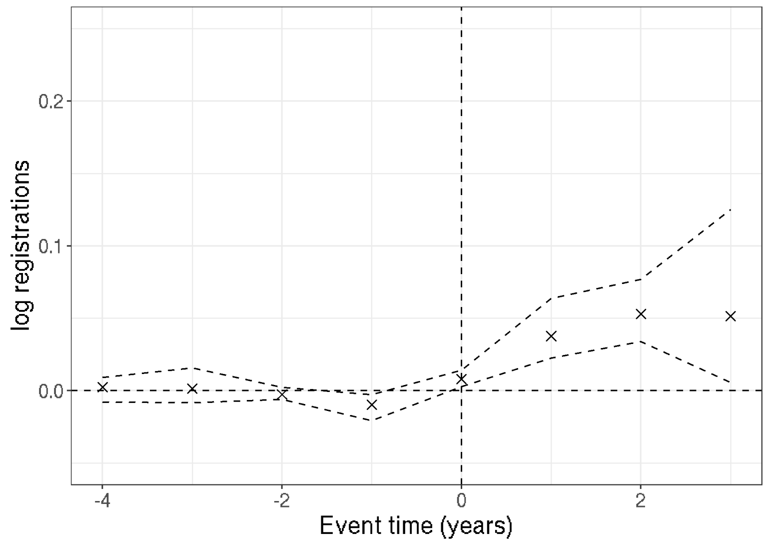

(a) Eligible vehicle registrations

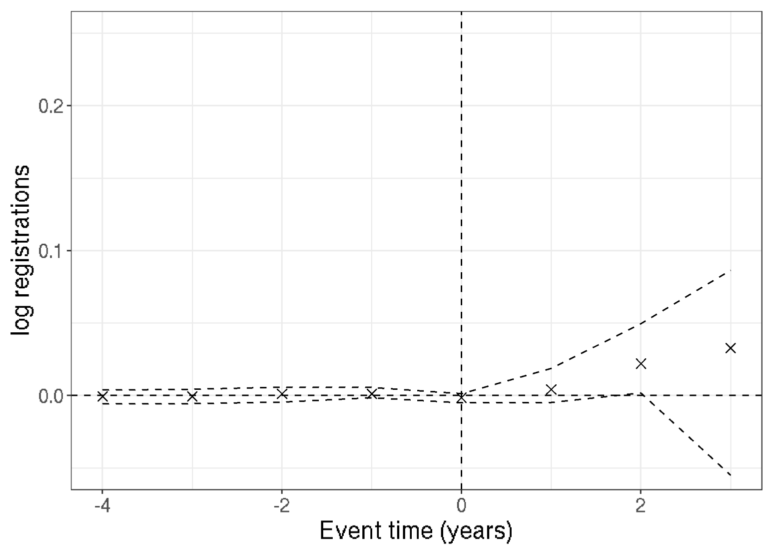

(b) Ineligible vehicle registrations

Table A7: DMV registrations and liens

Table A7 Panel A shows the the impact of ride share entry using DMV registration data between 2010 and 2016. The left-hand side variable is log new registrations. *Post* is an indicator for ride share entry. *Eligible* is an indicator for vehicle ride share eligibility. *Low income* is an indicator for whether the zip code is in the bottom 50% of zip codes in the CBSA by median income. All columns include zip-manufacturer-eligible and qtr-manufacturer-eligible fixed effects and separate linear time trends for income interacted with eligibility. Column (4) contains zip-quarter fixed effects. Panel B shows the impact of ride share entry on auto financing at the registration level. The left-hand side is an indicator for whether there is a lien. Columns (1)–(3) contain zip, time times low and high income fixed effects, and make-model fixed effects, with separate linear time trends for income interacted with eligibility. Column (4) has zip $\times$ time fixed effects. Standard errors in parentheses are clustered at the CBSA-quarter level.

Panel A: New auto registrations, select states, DMV data

	Log new registrations			
	(1)	(2)	(3)	(4)
Post	-0.0004	-0.005	0.007	-
	(0.016)	(0.017)	(0.014)	-
Post $\times$ Eligible	-	-	-0.018^{***}	-0.015^{**}
	-	-	(0.007)	(0.007)
Post $\times$ Low income	-	0.010^{*}	-0.003	-
	-	(0.006)	(0.005)	-
Post $\times$ Low income $\times$ Eligible	-	-	0.019^{***}	0.019^{***}
	-	-	(0.005)	(0.005)
Zip$\times$Manufacturer$\times$Eligible FE	Y	Y	Y	Y
Qtr$\times$Manufacturer$\times$Eligible FE	Y	Y	Y	Y
Zip$\times$Qtr FE	N	N	N	Y
Income$\times$Eligible time trends	Y	Y	Y	Y
Observations	3,348,566	3,348,566	3,348,566	3,348,566
R^2	0.454	0.454	0.454	0.472

Note: *p<0.1; **p<0.05; ***p<0.01

Panel B: Vehicle liens, select states, DMV data

	Has lien			
	(1)	(2)	(3)	(4)
Post	0.006^{**}	0.009	-0.011^{*}	-
	(0.002)	(0.006)	(0.006)	-
Post $\times$ Eligible	-	-	0.024^{***}	0.017^{***}
	-	-	(0.004)	(0.005)
Post $\times$ Low income	-	-0.016^{**}	-0.026^{***}	-
	-	(0.008)	(0.008)	-
Post $\times$ Low income $\times$ Eligible	-	-	0.018^{***}	0.013^{***}
	-	-	(0.004)	(0.004)
Qtr $\times$ Low income FE	Y	Y	Y	N
Zip FE	Y	Y	Y	N
Zip $\times$ Qtr FE	N	N	N	Y
Make-model FE	Y	Y	Y	Y
Income $\times$ Eligible time trends	Y	Y	Y	Y
Observations	1,709,002	1,132,906	1,132,906	1,132,906
R^2	0.252	0.295	0.073	0.308

Note: *p<0.1; **p<0.05; ***p<0.01

Table A8: Auto loan performance following ride share entry

Table A8 shows how auto loan performance changes in response to ride share entry. The specification exactly mirrors that in Regressions (1) and (2) with the exception of the left-hand side variable, which is the percentage of auto loans that are in serious delinquency (60+ days) within one year of origination:

$Performance_{zt} = \beta Post_{zt} + \gamma_t + \gamma_z + \epsilon_{zt} Performance_{zt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Low\ Income_z + \gamma_z + \gamma_{Income,t} + \epsilon_{zt}$

Data are from Equifax between 2010 and 2016. Standard errors, clustered at the CBSA-quarter level, are in parentheses.

	Dependent variable:		
	% In Default		
	(1)	(2)	(3)
Post	-0.0003^{***}	-0.001^{***}	-0.0005^{***}
	(0.0001)	(0.0001)	(0.0001)
Post × Low Wage		0.0005^{**}	-
	-	(0.0002)	-
Post × High transport share	-	-	0.0004^{*}
	-	-	(0.0002)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr × Low income FE	N	Y	N
Qtr × High transport FE	N	N	Y
Observations	244,156	244,156	244,156
R^2	0.590	0.593	0.591
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A9: Employment after ride share entry

Table A9 shows the result of regression (24) using IRS tax filing data between 2010 and 2016 broken out by business versus wage or salary income:

$$\log Filings_{bzt} = \beta Post_{zt} + FE + \epsilon_{bzt}$$

The regression is at the zip-year level. Only zips that eventually receive treatment are in the regression. The left-hand side variable is the log number of tax filings. *Post* is an indicator variable taking the value 1 after ride share has entered. All columns include year and zip fixed effects. Column (1) considers all filings. Column (2) considers filings with wage or salary income. Column (3) considers filings with business income. Standard errors, in parenthesis, are clustered at the CBSA-year level.

	Dependent variable:		
	log filings	log wage or salary filings	log business filings
	(1)	(2)	(3)
Post	0.003	0.003	0.013^{***}
	(0.002)	(0.002)	(0.004)
Zip FE	Y	Y	Y
Year FE	Y	Y	Y
Observations	72,646	72,646	72,646
R^2	0.999	0.999	0.989
Residual Std. Error (df = 58820)	0.041	0.040	0.165
Note:			*p<0.1; **p<0.05; ***p<0.01

Figure A10: Experimental design for bankruptcy flag removal

Figure A10 shows experimental design of the bankruptcy flag removal study. Between 8 and 12 years prior to ride share entry, all consumers in the study file for Chapter 7 bankruptcy. Between 1 and 2 years before ride share entry, the unconstrained group has their bankruptcy flags exogenously removed. They enter the event window, one year before and one year after ride share entry, with no bankruptcy filings on their credit reports. Between 1 and 2 years after ride share entry, the constrained group has their bankruptcy flags exogenously removed. For the entirety of the event window, their bankruptcy flag is present. The experiment compares the borrowing response to ride share entry of the constrained borrowers to the unconstrained borrowers.

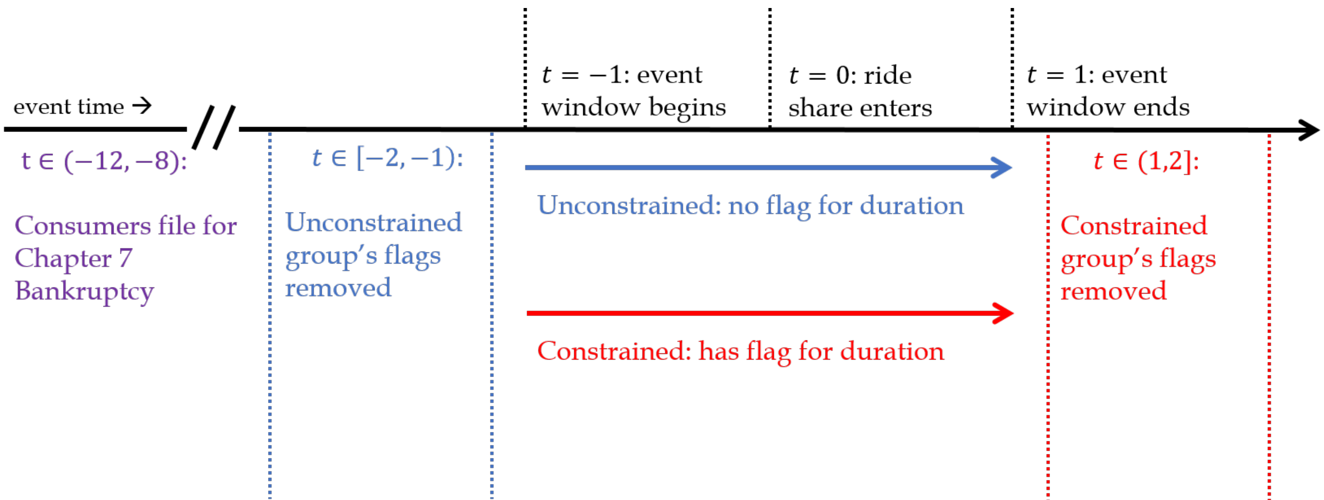

Figure A11: Financial constraints and auto loans after ride share entry

Figure A11 shows the impact of financial constraints on auto loan origination following ride share entry. The left-hand side variable is log auto loan originations, and the sample is split according to financial constraints as measured at the zip code level. Panels (a) and (b) split the sample by 2010 bank share of auto lending, with Panel (a) showing low-bank (unconstrained) zip codes and panel (b) showing high-bank (constrained) zip codes. Panels (c) and (d) split the sample by 2010 consumer finance defaults, with Panel (c) showing low-default (unconstrained) zip codes and panel (d) showing high-default (constrained) zip codes. Data are from Equifax. Gray bars show 95% confidence intervals.

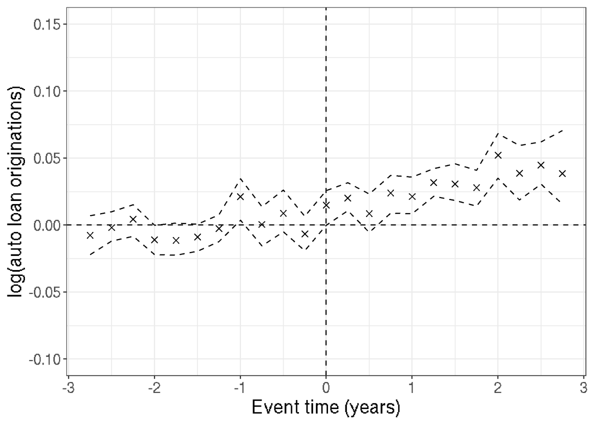

(a) Vehicle sales in low-bank zip codes

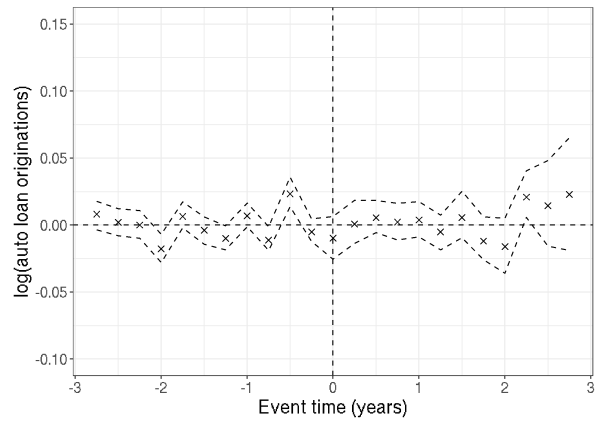

(b) Vehicle sales in high-bank zip codes

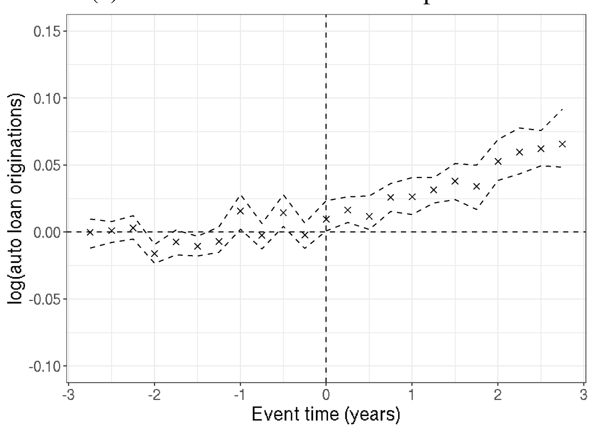

(c) Vehicle sales in low-default zip codes

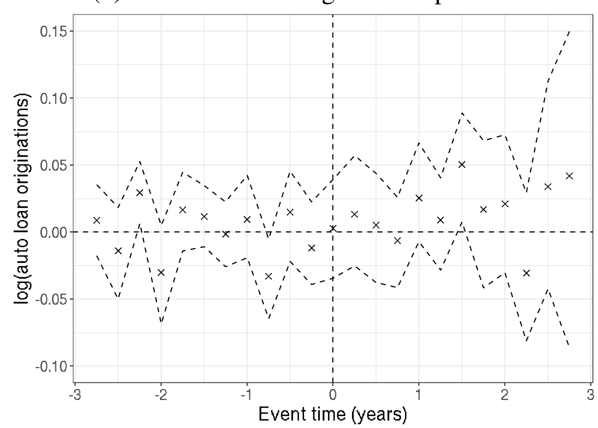

(d) Vehicle sales in high-default zip codes

Figure A12: Financial constraints and auto sales after ride share entry

Figure A12 shows the impact of financial constraints on new auto sales following ride share entry. The left-hand side variable is log new auto sales, and the sample is split according to financial constraints as measured at the zip code level. Panels (a) and (b) split the sample by 2010 bank share of auto lending, with Panel (a) showing low-bank (unconstrained) zip codes and panel (b) showing high-bank (constrained) zip codes. Panels (c) and (d) split the sample by 2010 consumer finance defaults, with Panel (c) showing low-default (unconstrained) zip codes and panel (d) showing high-default (constrained) zip codes. Data are from RL Polk. Gray bars show 95% confidence intervals.

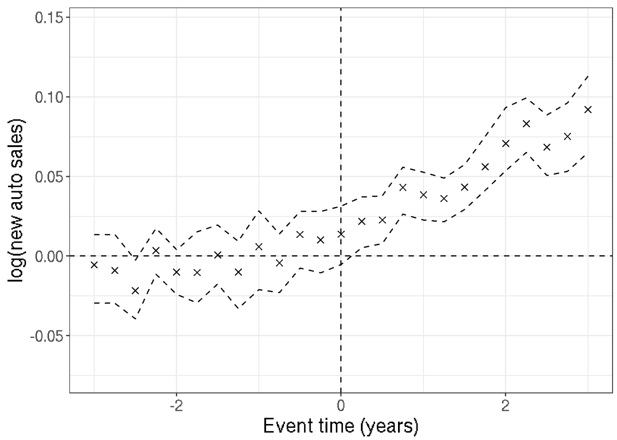

(a) Vehicle sales in low-bank zip codes

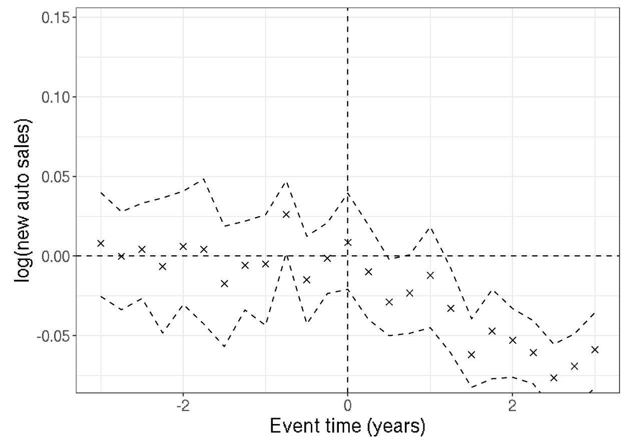

(b) Vehicle sales in high-bank zip codes

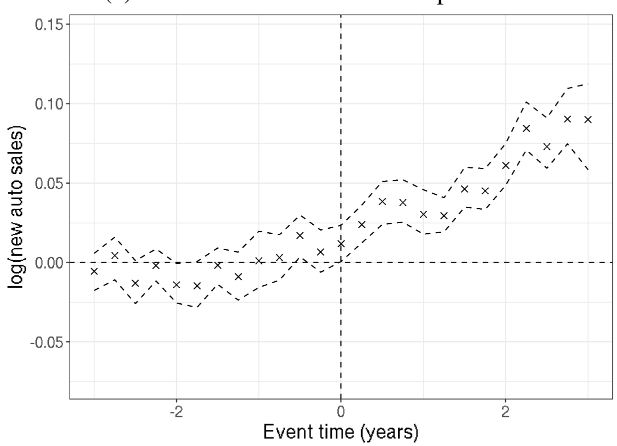

(c) Vehicle sales in low-default zip codes

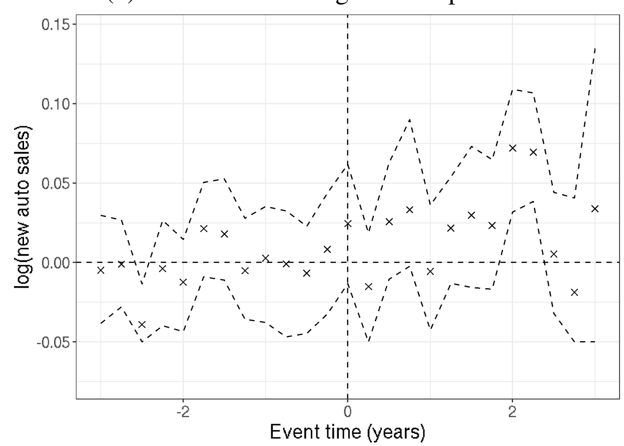

(d) Vehicle sales in high-default zip codes

Figure A13: Financial constraints and tax filings after ride share entry

Figure A13 shows the impact of financial constraints on employment following ride share entry. The left-hand side variable is log tax filings, and the sample is split according to financial constraints as measured at the zip code level. Panels (a) and (b) split the sample by 2010 bank share of auto lending, with Panel (a) showing low-bank (unconstrained) zip codes and panel (b) showing high-bank (constrained) zip codes. Panels (c) and (d) split the sample by 2010 consumer finance defaults, with Panel (c) showing low-default (unconstrained) zip codes and panel (d) showing high-default (constrained) zip codes. Data are from the IRS Summary of Income. Gray bars show 95% confidence intervals.

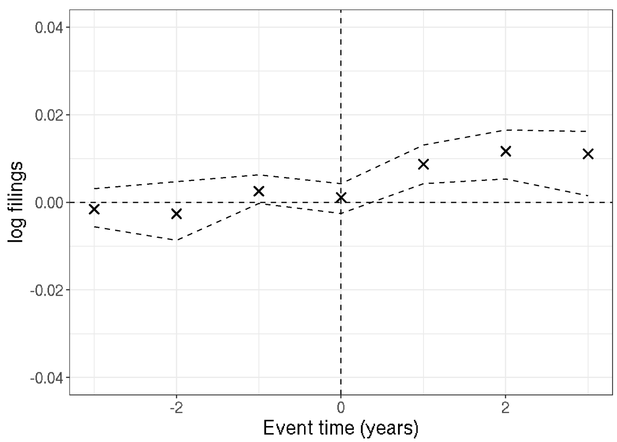

(a) Tax filings in low-bank zip codes

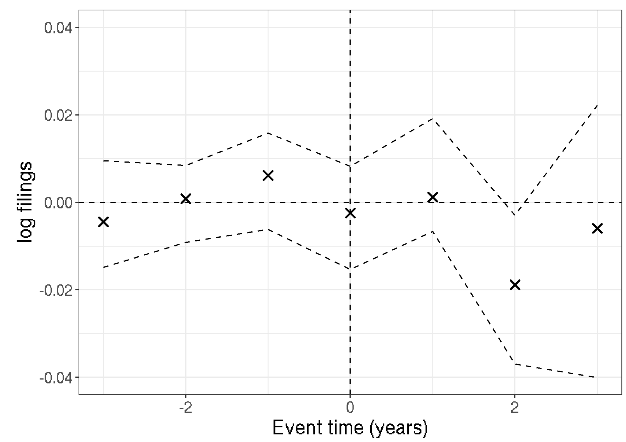

(b) Tax filings in high-bank zip codes

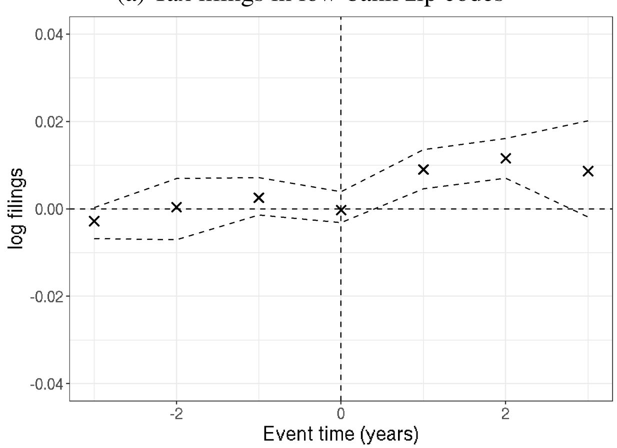

(c) Tax filings in low-default zip codes

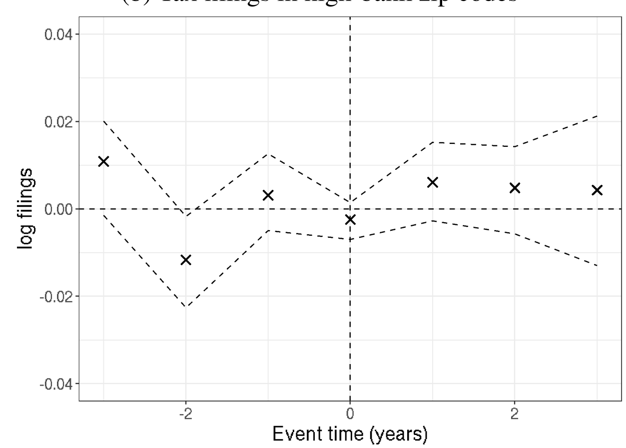

(d) Tax filings in high-default zip codes

Table A14: Financial constraints and ride share entry

Table A14 shows the impact of credit supply shocks on auto loan origination (Panel A) and auto sales and employment (Panel B) following ride share entry. The regressions are on the quarter-zip level for loans and sales and year-zip level for tax filings. *High bank share* is an indicator for whether the bank share of auto lending in the zip code in 2010 was above the 90th percentile in the CBSA. *High default share* is an indicator for whether the percentage of consumer loans in default in the zip code in 2010 was above the 90th percentile. All columns include *Post* × *High income* controls, not shown in this table. Additionally, all columns include zip fixed effects, and *High bank* or *High default* × *Low income* × time fixed effects. Data are from Equifax, RL Polk, and the IRS SOI between 2010 and 2016. Standard errors, in parentheses, are clustered at the CBSA-quarter level.

Panel A: Financial constraints and auto lending following ride share entry

	Dependent variable:	
	log new originations	
	(1)	(2)
Post	0.021***	0.024***
	(0.005)	(0.005)
Post × High bank share	−0.080***	-
	(0.017)	-
Post × High default share	-	−0.078***
	-	(0.019)
Zip FE	Y	Y
Credit×Income× Time FE	Y	Y
Post×Income	Y	Y
Observations	320,279	320,279
R^2	0.978	0.978
Note:	*p<0.1; **p<0.05; ***p<0.01	

Panel B: Financial constraints and the real effects of ride share entry

	Dependent variable:			
	log sales		log tax filings	
	(1)	(2)	(3)	(4)
Post	0.022***	0.025***	0.009***	0.009***
	(0.006)	(0.006)	(0.002)	(0.002)
Post × High bank share	−0.038**	-	−0.008**	-
	(0.017)	-	(0.004)	-
Post × High default share	-	−0.048***	-	−0.010**
	-	(0.014)	-	(0.004)
Zip FE	Y	Y	Y	Y
Credit×Income× Time FE	Y	Y	Y	Y
Post×Income	Y	Y	Y	Y
Observations	454,949	454,949	83,793	83,793
R^2	0.969	0.969	0.999	0.999
Note:	*p<0.1; **p<0.05; ***p<0.01			

Figure A15: Estimation diagnostics

Figure A15 shows actual (blue sold line and 'x' markers) and model generated (red dashed line and 'o' markers) vehicle financing rates versus individual income. Panel (a) shows the results over the whole sample. Panel (b) shows the results pre ride-share entry. Panel (c) shows the results post ride-share. Panel (d) shows actual versus predicted financing. Gray bars represent 95% confidence intervals.

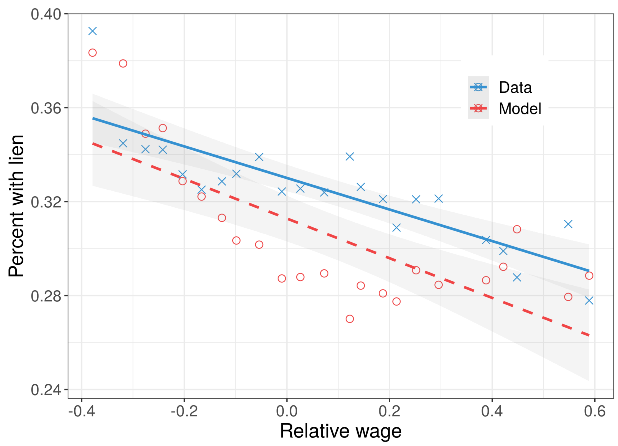

(a) All periods

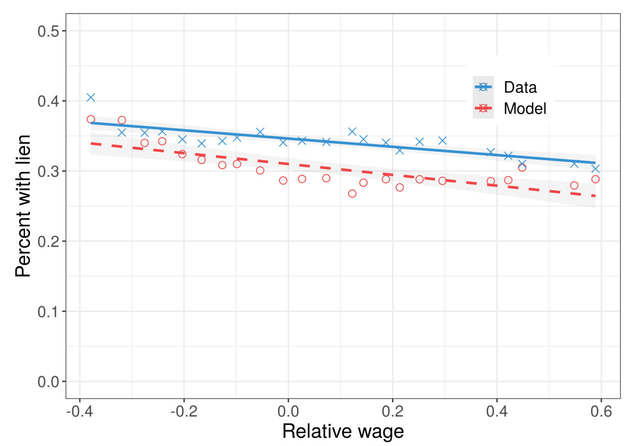

(b) Pre ride-share entry

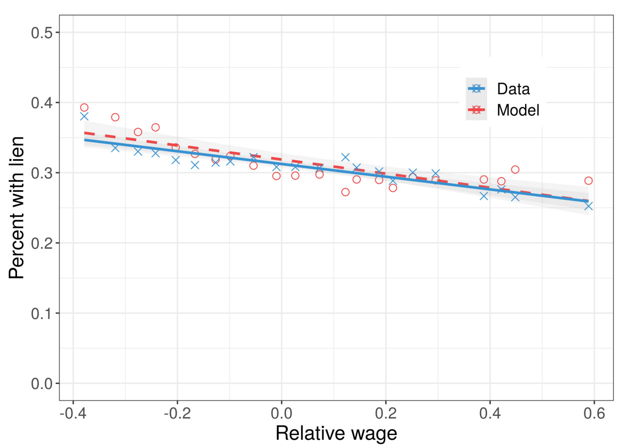

(c) Post ride-share entry

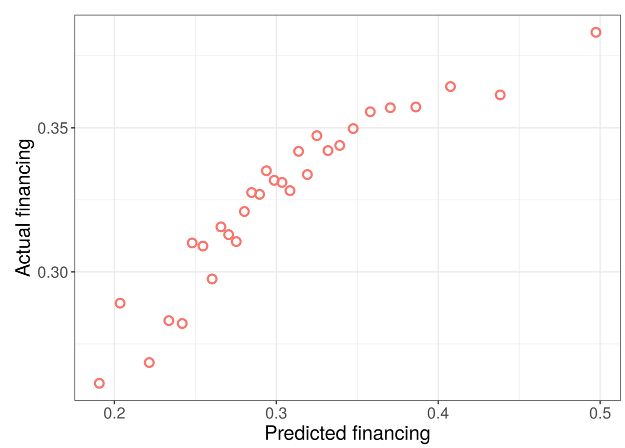

(d) Actual versus model-predicted financing

Figure A16: Counterfactual costs of finance, elasticity sensitivity

Figure A16 shows the impact of increasing financing costs for low (0.42), baseline (0.57), and high (0.72) levels of demand elasticity. In each panel, the x-axis shows financing costs relative to the baseline estimated value capitalized value of $14,435, and the vertical dashed line indicates to that value. Panel (a) shows the quantity of ride share rides, normalized to the current value. Panel (b) shows the hourly price of ride share rides. Panel (c) shows drivers' hourly net income, defined as yearly income from driving for ride share full time less vehicle financing costs. Panel (d) shows the percentage of drivers using financing. Panel (e) shows individual driver welfare. Panel (f) shows driver plus rider welfare.

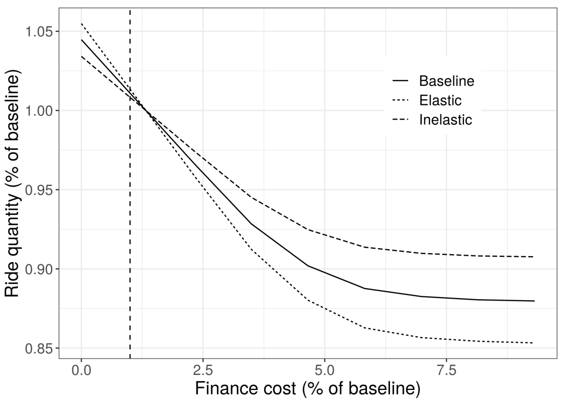

(a) Ride share ride quantities

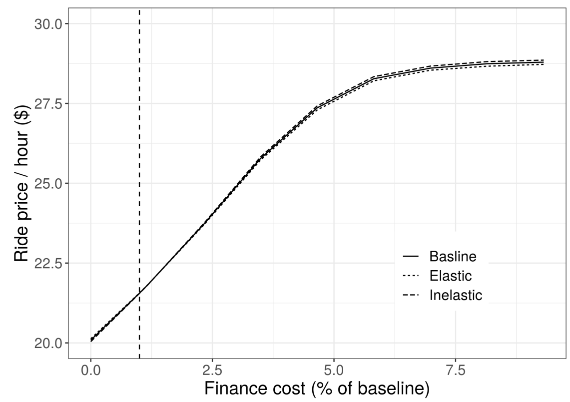

(b) Ride share ride prices

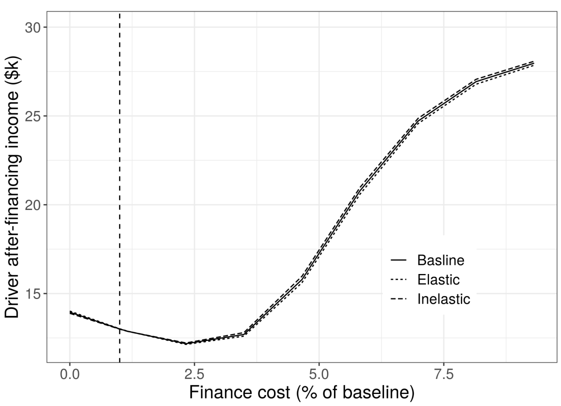

(c) Driver net income

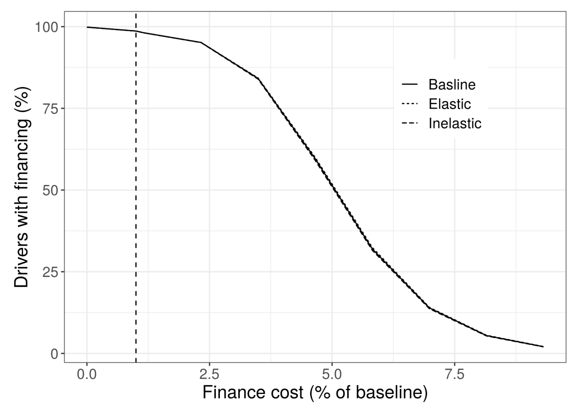

(d) Driver outside income opportunities

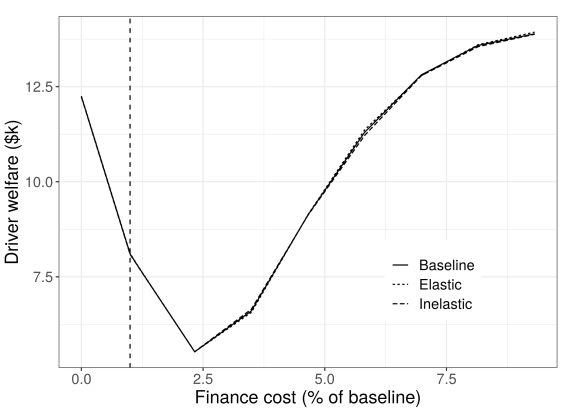

(e) Individual driver welfare

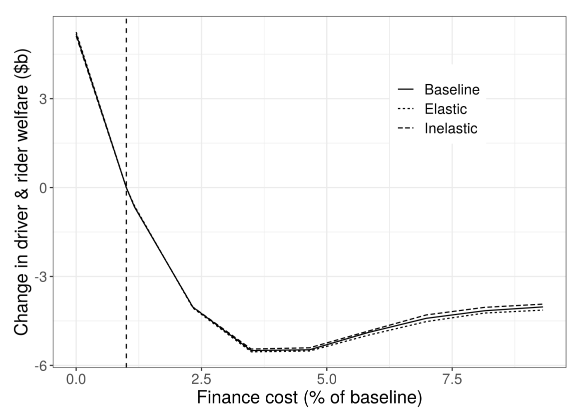

(f) Driver plus rider welfare

Figure A17: Counterfactual systems of capital allocation, elasticity sensitivity

Figure A17 compares systems of capital allocation for low (0.42), baseline (0.57), and high (0.72) levels of demand elasticity. The traditional firm scenario envisions firms owning financed cars and hiring drivers at minimum wage. The gig economy scenario is the current structure, where drivers work as independent contractors and own their cars. The car share scenario envisions individuals owning cars and renting them frictionlessly to drivers earning minimum wage. Panel (a) shows ride quantities relative to the baseline. Panel (b) shows ride prices per hour. Panel (c) shows driver net income. Panel (d) shows the percentage of drivers using financing. Panel (e) shows driver plus rider welfare.

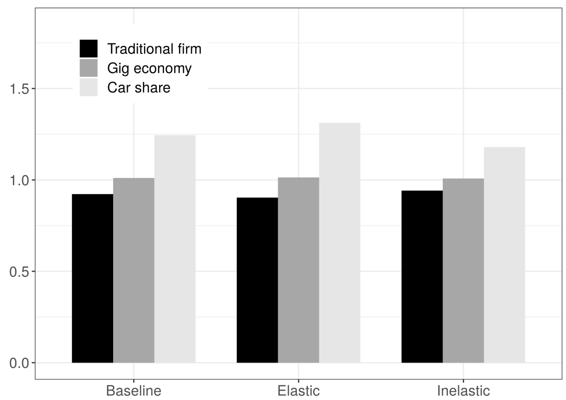

(a) Ride share ride quantities

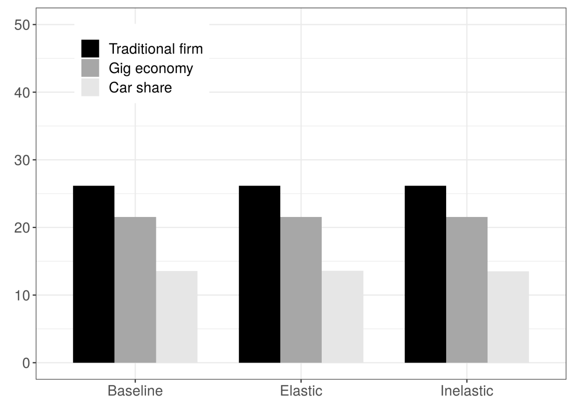

(b) Ride share ride prices

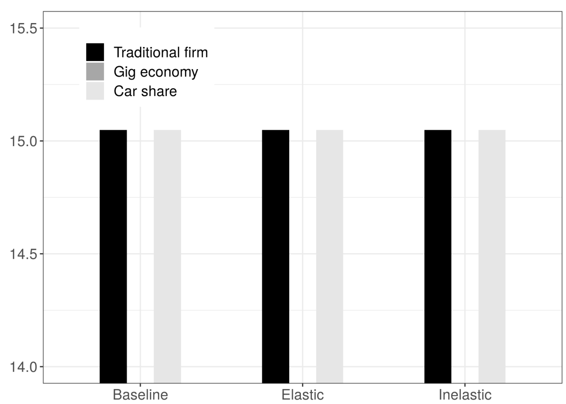

(c) Driver net income

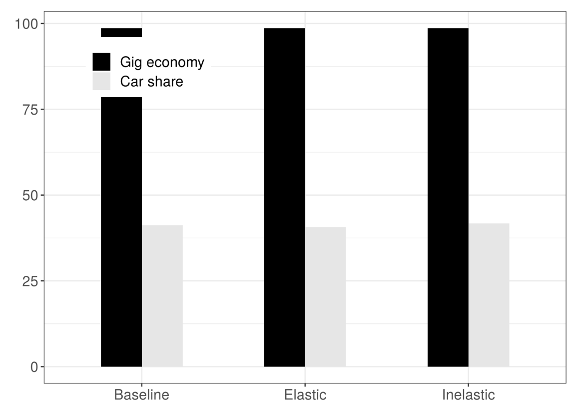

(d) Vehicle owner outside income

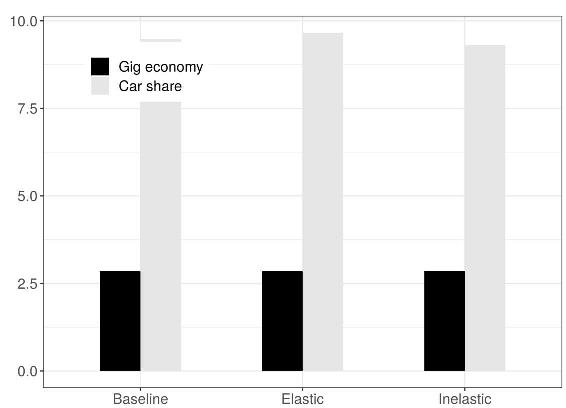

(e) Driver and rider welfare

Figure A18: Counterfactual costs of finance, commission sensitivity

Figure A18 shows the impact of increasing financing costs for low (15%), baseline (25%), and high (35%) ride share commissions. In each panel, the x-axis shows financing costs relative to the baseline estimated value capitalized value of $14,435, and the vertical dashed line indicates to that value. Panel (a) shows the quantity of ride share rides, normalized to the current value. Panel (b) shows the hourly price of ride share rides. Panel (c) shows drivers' hourly net income, defined as yearly income from driving for ride share full time less vehicle financing costs. Panel (d) shows the percentage of drivers using financing. Panel (e) shows individual driver welfare. Panel (f) shows driver plus rider welfare.

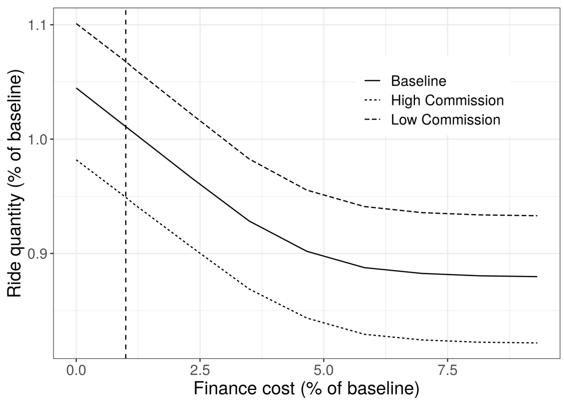

(a) Ride share ride quantities

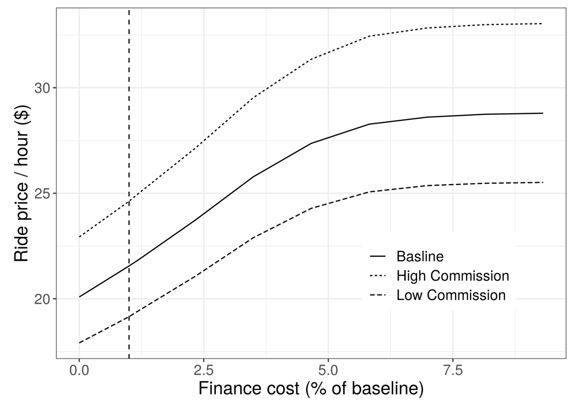

(b) Ride share ride prices

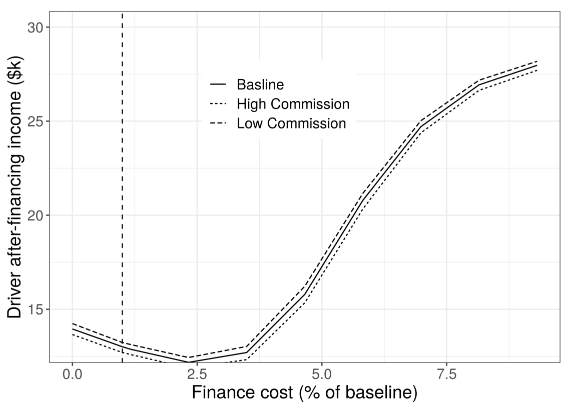

(c) Driver net income

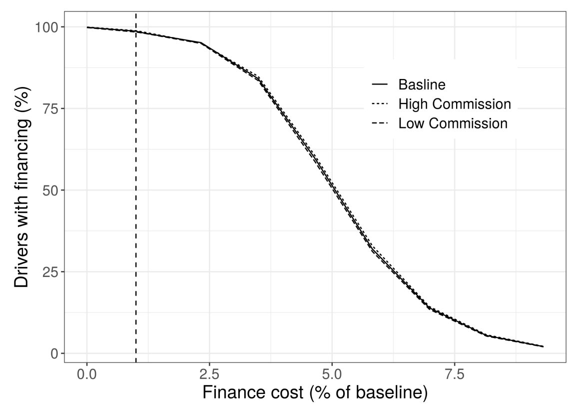

(d) Driver outside income opportunities

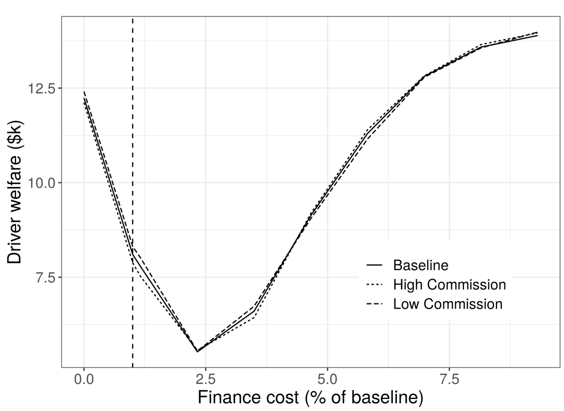

(e) Individual driver welfare

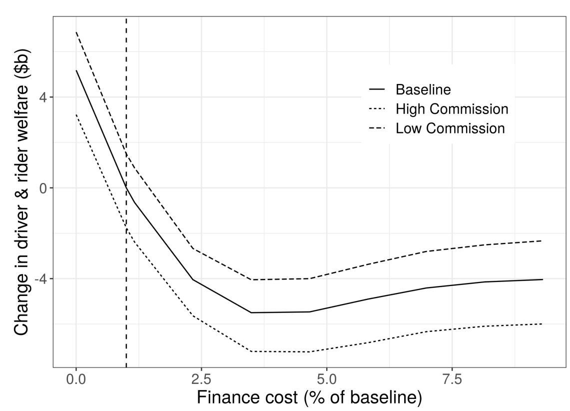

(f) Driver and rider welfare

Figure A19: Counterfactual systems of capital allocation, elasticity sensitivity

Figure A19 compares systems of capital allocation for low (15%), baseline (25%), and high (35%) ride share commissions. The traditional firm scenario envisions firms owning financed cars and hiring drivers at minimum wage. The gig economy scenario is the current structure, where drivers work as independent contractors and own their cars. The car share scenario envisions individuals owning cars and renting them frictionlessly to drivers earning minimum wage. Panel (a) shows ride quantities relative to the baseline. Panel (b) shows ride prices per hour. Panel (c) shows driver net income. Panel (d) shows the percentage of drivers using financing. Panel (e) shows driver plus rider welfare.

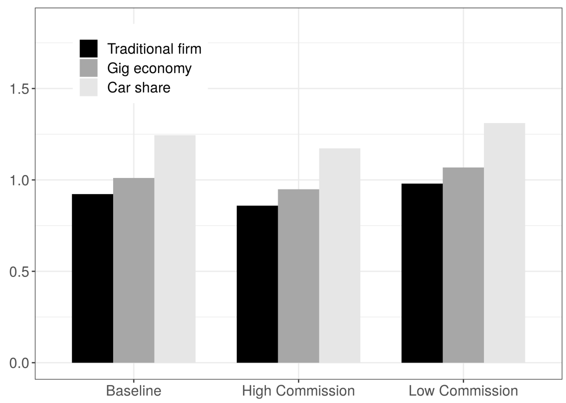

(a) Ride share ride quantities

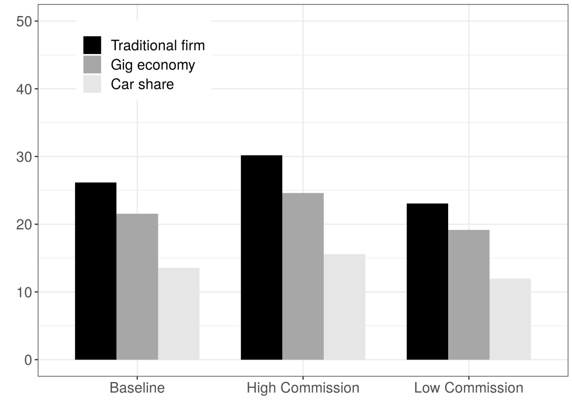

(b) Ride share ride prices

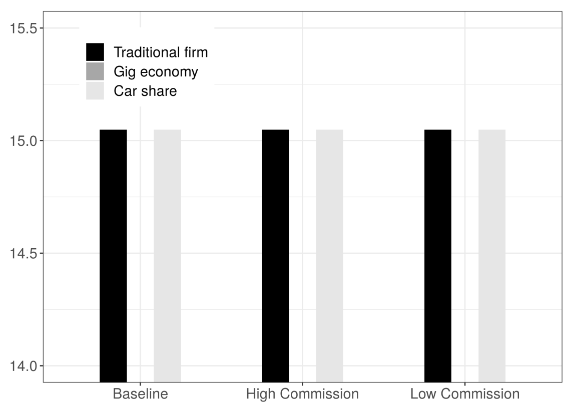

(c) Driver net income

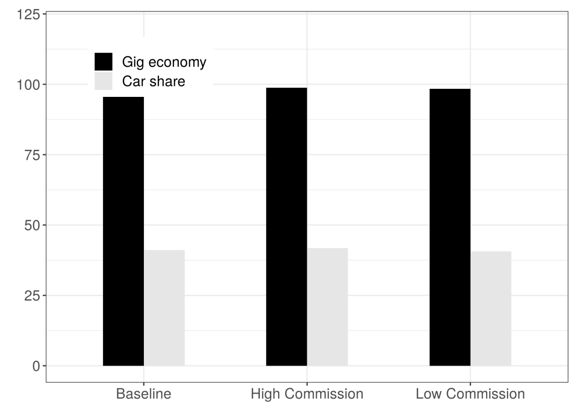

(d) Vehicle owner outside income

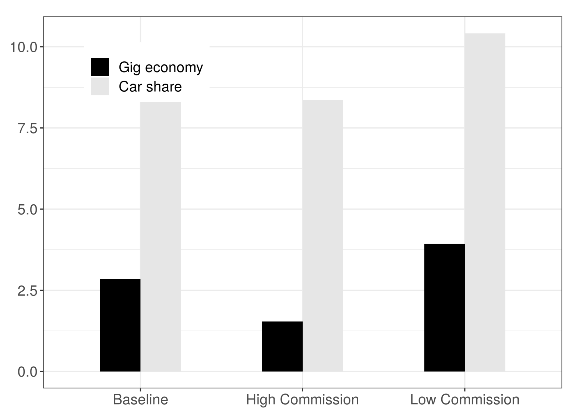

(e) Driver and rider welfare

Table A20: Nationwide auto sales and loans after ride share entry, placebo test

Table A20 shows the placebo test corresponding to Table 1, where entry location is randomized. Panel A is the nationwide RL Polk data; Panel B is the nationwide Equifax data. See Table 1 for detailed specification and variable descriptions.

Panel A: New auto sales, nationwide, RL Polk data

	Dependent variable:		
	Log sales		
	(1)	(2)	(3)
Post	0.0004	0.001	0.001
	(0.002)	(0.002)	(0.003)
Post × Low income	-	−0.002	-
	-	(0.004)	-
Post × High transport share	-	-	−0.001
	-	-	(0.004)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr × Low income FE	N	Y	N
Qtr × High transport FE	N	N	Y
Observations	299,332	299,332	299,332
R^2	0.967	0.967	0.967
Note:	*p<0.1; **p<0.05; ***p<0.01		

Panel B: New auto loan originations, nationwide, Equifax data

	Dependent variable:		
	Log new originations		
	(1)	(2)	(3)
Post	−0.002	−0.001	−0.004
	(0.002)	(0.002)	(0.002)
Post × Low income	-	−0.002	-
	-	(0.003)	-
Post × Transport share	-	-	0.003
	-	-	(0.003)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr × Low income FE	N	Y	N
Qtr × High transport FE	N	N	Y
Observations	299,332	299,332	299,332
R^2	0.974	0.974	0.974
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table A21: DMV registrations and liens, placebo test

Table A21 shows the placebo test corresponding to Table A7, where entry location is randomized. Panel A is the nationwide RL Polk data; Panel B is the select state results using DMV data. See TTable A7 for detailed specification and variable descriptions.

Panel A: New auto registrations, select states, DMV data

	Dependent variable:			
	Log new registrations			
	(1)	(2)	(3)	(4)
Post	-0.002	0.003	0.0004	-
	(0.003)	(0.003)	(0.003)	-
Post × Eligible	-	-	0.004	0.003
	-	-	(0.003)	(0.003)
Post × Low income	-	-0.010^{**}	-0.008^{**}	-
	-	(0.004)	(0.004)	-
Post × Low income × Eligible	-	-	-0.002	-0.003
	-	-	(0.003)	(0.003)
Zip×Manufacturer×Eligible FE	Y	Y	Y	Y
Qtr×Manufacturer×Eligible FE	Y	Y	Y	Y
Zip×Qtr FE	N	N	N	Y
Income×Eligible time trends	Y	Y	Y	Y
Observations	3,348,566	3,348,566	3,348,566	3,348,566
R^2	0.454	0.454	0.454	0.472
Note:				*p<0.1; **p<0.05; ***p<0.01

Panel B: Vehicle liens, select states, DMV data

	Dependent variable:			
	Has lien			
	(1)	(2)	(3)	(4)
Post	0.004	0.008	-0.009	-
	(0.002)	(0.005)	(0.006)	-
Post × Eligible	-	-	0.023^{***}	0.024^{***}
	-	-	(0.004)	(0.005)
Post × Low income	-	-0.008	-0.021^{***}	-
	-	(0.007)	(0.008)	-
Post × Low income × eligible	-	-	0.020^{***}	0.012^{***}
	-	-	(0.004)	(0.005)
Qtr × Low income FE	Y	Y	Y	N
Zip FE	Y	Y	Y	N
Zip × Qtr FE	N	N	N	Y
Make-model FE	Y	Y	Y	Y
Eligible time trends	Y	Y	Y	Y
Observations	1,710,358	1,134,335	1,134,335	1,134,335
R^2	0.252	0.295	0.073	0.308
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A22: Employment after ride share entry, placebo test

Table A22 shows the placebo analysis corresponding to Table (2) Panel A using IRS tax filing data between 2010 and 2016. The regression is at the zip-year-tax bucket level. Only zips that eventually receive treatment are in the regression. The left-hand side variable is the log number of tax filings. *Post* is an indicator variable taking the value 1 after ride share has entered, but treatment has been randomly assigned for this placebo test. *AGI<25k* is an indicator taking the value 1 for the AGI bucket below $25,000. All columns include tax-bucket times year and tax-bucket times zip fixed effects. Column (3) additionally contains zip times year fixed effects. Standard errors, in parenthesis, are clustered at the CBSA-year level.

	Dependent variable:		
	log(n.filings)		
	(1)	(2)	(3)
Post	0.0002	−0.0001	
	(0.001)	(0.002)	(0.000)
Post × (AGI<25k)	-	0.001	0.001
	-	(0.002)	(0.002)
(AGI<25k)×Year FE	Y	Y	Y
(AGI<25k)×Zip FE	Y	Y	Y
Zip×Year FE	N	N	Y
Observations	152,215	152,215	152,215
R^2	0.995	0.995	0.997
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$		

Table A23: Vehicle utilization after ride share entry, placebo test

Table A23 shows the placebo test corresponding to Table 2 Panel B, where entry location is randomized. See Table 2 for detailed specification and variable descriptions.

	Dependent variable:					
	Miles / year (k)					
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.127	0.198	0.461	-	0.739	-
	(0.128)	(0.313)	(0.321)	-	(0.745)	-
Post × Low income	-	−0.077	-	-	−0.308	-
	-	(0.315)	-	-	(0.754)	-
Post × Eligible	-	-	−0.381	−0.256	−0.629	−0.801
	-	-	(0.344)	(0.394)	(0.782)	(0.833)
Post × Low income × Eligible	-	-	-	-	0.276	0.598
	-	-	-	-	(0.791)	(0.850)
Zip × Eligible FE	Y	Y	Y	Y	Y	Y
Qtr × Eligible FE	Y	Y	Y	Y	Y	Y
Zip × Qtr FE	N	N	N	Y	N	Y
Observations	122,674	122,645	122,674	122,674	122,645	122,645
R^2	0.048	0.048	0.048	0.178	0.048	0.178

Note: *p<0.1; **p<0.05; ***p<0.01

Table A24: Zip-level credit constraints, placebo test

Table A24 shows the placebo test corresponding to Table A14, where entry location is randomized. See Table A14 Panels A and B for detailed specification and variable descriptions.

	Dependent variable:					
	log new originations		log sales		log tax filings	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	−0.002	−0.003	0.003	0.003	−0.0002	−0.0004
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Post × High bank share	0.005	-	−0.004	-	−0.001	-
	(0.008)	-	(0.009)	-	(0.002)	-
Post × High default share	-	0.014*	-	0.001	-	−0.0003
	-	(0.008)	-	(0.009)	-	(0.002)
Zip FE	Y	Y	Y	Y	Y	Y
Credit×Wage× Time FE	Y	Y	Y	Y	Y	Y
Post×Wage	Y	Y	Y	Y	Y	Y
Observations	321,554	321,554	456,386	456,386	82,386	82,386
R^2	0.973	0.973	0.964	0.964	0.999	0.999
Note:						*p<0.1; **p<0.05; ***p<0.01