

THE UNIVERSITY OF CHICAGO

FINANCING THE GIG ECONOMY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
AND
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
GREGORY EDWARD BUCHAK

CHICAGO, ILLINOIS

JUNE 2019

CONTENTS

LIST OF FIGURES	iv
LIST OF TABLES	v
ACKNOWLEDGMENTS	vi
ABSTRACT	vii
1 INTRODUCTION	1
2 DATA AND INSTITUTIONAL BACKGROUND	8
2.1 Data	8
2.2 Institutional Background	8
3 RIDE SHARE'S REAL EFFECTS AND THE ROLE OF FINANCE	10
3.1 Who owns cars before ride share?	10
3.2 What happens when ride share enters?	11
3.2.1 Auto sales and loans	11
3.2.2 Employment	14
3.2.3 Vehicle utilization	16
3.3 Financial constraints and ride share growth	18
3.3.1 Borrower-level evidence from bankruptcy flag removal	19
3.3.2 Zip-level evidence	21
4 A STRUCTURAL MODEL OF RIDE SHARE CAPITAL	24
4.1 Model description	25
4.1.1 Individual consumer's problem	26
4.1.2 Consumer preferences and aggregation	33
4.1.3 Demand for ride share services	34
4.1.4 Estimation	35
4.2 Counterfactuals	37
4.2.1 Financing costs in the gig economy	38
4.2.2 Capital allocation alternatives to the gig economy	40
4.2.3 Driver quantity restrictions	43
5 DISCUSSION AND CONCLUSION	45
5.1 Related literature	45
5.1.1 Consumer finance	45
5.1.2 Corporate finance and productivity	46
5.1.3 Uber, Lyft, and ride share	47
5.1.4 Methodology	48
5.2 Discussion	48
A FIGURES	51

B	TABLES	68
C	SUPPORTING ANALYSES	75
	C.1 Data	76
	C.2 The timing and location of ride share entry	81
	C.3 Ex-ante determinants of vehicle ownership	84
	C.4 DMV Registrations and Liens	86
	C.5 Loan performance after ride share entry	93
	C.6 Model sensitivity to demand assumptions	95
	C.7 Placebo tests	101
	BIBLIOGRAPHY	106

LIST OF FIGURES

A.1	Timing of ride share entry	52
A.2	Vehicle ownership by household income prior to ride share entry	53
A.3	Vehicle sales and auto loans after ride share entry	54
A.4	Vehicle sales and auto loans after ride share entry by zip code income	55
A.5	Tax filings following ride share entry	56
A.6	Vehicle utilization after ride share entry	57
A.7	Bankruptcy disclosure and ride share financing	58
A.8	Experimental design for bankruptcy flag removal	59
A.9	Financial constraints and auto loans after ride share entry	60
A.10	Financial constraints and auto sales after ride share entry	61
A.11	Financial constraints and tax filings after ride share entry	62
A.12	Estimation diagnostics: financing and wage	63
A.13	Counterfactual costs of finance in the gig economy, prices, quantities, and wages	64
A.14	Counterfactual systems of capital allocation	65
A.15	Finance costs and counterfactual systems of capital allocation	66
A.16	Counterfactual limits on drivers	67
C.1	New vehicle registrations following ride share entry	91
C.2	Counterfactual costs of finance, elasticity sensitivity	97
C.3	Counterfactual systems of capital allocation, elasticity sensitivity	98
C.4	Counterfactual costs of finance, commission sensitivity	99
C.5	Counterfactual systems of capital allocation, elasticity sensitivity	100

LIST OF TABLES

B.1	Nationwide auto sales and loan originations after ride share entry	69
B.2	Employment after ride share entry	70
B.3	Vehicle utilization after ride share entry	71
B.4	Bankruptcy disclosure and ride share financing	72
B.5	Financial constraints and ride share entry	73
B.6	Model parameters	74
C.1	Summary statistics	79
C.2	New registrations and new sales comparison	80
C.3	Timing of ride share entry	83
C.4	Vehicle stocks and sales before ride share entry	85
C.5	DMV registrations and liens	92
C.6	Auto loan performance following ride share entry	94
C.7	Nationwide auto sales and loans after ride share entry, placebo test	101
C.8	DMV registrations and liens, placebo test	102
C.9	Employment after ride share entry, placebo test	103
C.10	Vehicle utilization after ride share entry, placebo test	104
C.11	Zip-level credit constraints, placebo test	105

ACKNOWLEDGMENTS

I am grateful to my committee, Amit Seru, Amir Sufi, Gregor Matvos, Steve Kaplan, and Ufuk Akcigit. 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. TransUnion (the data provider) has the right to review the research before dissemination to ensure it accurately describes TransUnion data, does not disclose confidential information, and does not contain material it deems to be misleading or false regarding TransUnion, TransUnions partners, affiliates or customer base, or the consumer lending industry. TransUnion can require disputed parts of the research to be redacted.

ABSTRACT

I study the impact of credit constraints on gig economy penetration, capital allocation, and employment through the lens of Uber and Lyft. The low-income individuals for whom ride share driving is attractive often require financing to obtain cars. Exploiting the staggered entry of ride share across cities and within-city variation in income, I find that ride share entry coincides with sharp increases in auto loans, auto sales, employment, and vehicle utilization among low-income individuals. Within zip codes, these effects are concentrated among ride-share eligible vehicles. Using the exogenous removal of bankruptcy disclosures as a shock to credit availability, I find that financial constraints dampen these effects. Motivated by these facts, I build a structural model linking consumers' vehicle acquisition and utilization, ride share driving, and financing decisions. I quantify the distributional and welfare implications of changes in credit supply and market structure. An increase in financing costs forces finance-dependent low-income drivers from the market and replaces them with wealthier, less financially constrained drivers with significantly higher outside earning opportunities. After-interest driver income nearly doubles, but finance-dependent drivers do not capture these benefits. Introducing a frictionless rental market for cars allows drivers to utilize idle capital, increases ride quantities by 20%, and decreases ride prices by 35%. Organizing the industry around taxi companies that own cars has the opposite effect. Proposed policies to restrict the number of drivers impose significant welfare costs that fall primarily on riders rather than drivers. These results suggest that finance critically shapes the size and boundaries of the gig economy.

CHAPTER 1

INTRODUCTION

Disruptive technologies like Uber, Lyft, and Airbnb have opened up a wealth of new labor opportunities. Individuals with cars can drive in their spare time or deliver takeout, and households with extra space can house guests in their spare rooms. Facing new and flexible working opportunities, individuals and households have joined the so-called gig economy as drivers, delivery people, and hosts, catalyzing explosive growth. In 2016, there were nearly 800,000 active Uber drivers in the United States, and as of 2015, more than 500,000 American households listed rooms on Airbnb.

The gig economy's structure implies a close connection between its rapid growth and the allocation and utilization of physical capital. After all, one cannot be an Uber driver without an Uber car. These apps, as well as policy makers and the popular press, often portray ride share drivers as part-time workers who already owned cars before ride share entered the market.¹ This message misses an important aspect of ride share employment: ride share driving pays a relatively low wage,² and prior to ride share entry, cars were mostly in the hands of high- rather than low-income individuals. Ride share technologies thus contend with a mismatch between the individuals who benefit from ride share employment and those who own cars. Low-income individuals that want to become ride share drivers first need ride share eligible cars. The difficulty is that these low-income individuals are the least likely to have the financial capital necessary to buy cars outright, and therefore, many must borrow the necessary capital.

1. On the driver sign-up page, Uber promises, "Drive when you want," "Earn what you need," and "Don't wait to start making great money with your car." <https://www.uber.com/a/join-now/>. The Washington Post evaluates the hidden costs of driving for Uber, and says "Despite all of this, you may still decide that you would like to be your own boss and make extra cash in your spare time." <https://www.washingtonpost.com/news/get-there/wp/2015/02/20/the-hidden-costs-of-being-an-uber-driver/>. In discussing congestion following Uber entry, the New York Times reports on some transit officials who "view the apps as an ally in their efforts to persuade urban dwellers to resist car ownership," and Lyft states that "Lyft's goal is to reduce private car ownership." <https://www.nytimes.com/2017/03/06/nyregion/uber-ride-hailing-new-york-transportation.html>. Sites accessed on October 10, 2018.

2. Mishel (2018) estimates the average after-expense income is roughly \$11.77 per hour.

In this paper, I study capital allocation in the gig economy through the lens of Uber and Lyft. I ask whether the financial system effectively puts cars into the hands of low-income individuals seeking to become ride share drivers. I first provide evidence that ride share entry leads to sharp increases in auto sales, employment, and vehicle utilization. These increases are concentrated among low-income households and ride share-eligible vehicles. Crucially, these increases rely on financing. Auto lending spikes when ride share enters a market, and the real effects are smaller when financing is unavailable. Past bankruptcy disclosures on consumer credit reports, past consumer loan defaults, and a restricted supply of auto lending reduce consumer borrowing following ride share entry. Less borrowing leads to fewer auto sales and prevents households from joining ride share employment.

Given the importance of finance, I build a structural model to assess the equilibrium distributional and welfare consequences of costlier financing. Additionally, I model alternative mechanisms for allocating capital and the effects of proposed policies that limit ride share driver quantities. Costlier financing dramatically alters the driver composition by forcing low-income, finance-dependent individuals out of the market and replacing them with wealthier individuals not reliant on finance. These new drivers have roughly 50% higher non-ride share earnings potential. Ride share wages nearly double, even after accounting for increased financing costs, because supply reductions increase ride prices.

So long as finance is available, however, the gig economy's capital allocation mechanism, where independent drivers own their own cars, offers significant welfare benefits over a traditional taxi company that owns cars and hires labor and leads to 25% more rides at 33% lower prices. This occurs because cars act as both productive capital assets and convenient consumer durables. In consequence, lower ride share returns justify vehicle purchases. Some emerging technologies are now attempting to introduce short-term, on-demand rental markets connecting car-owning households with drivers.³ I find that such a market would lead to a further 22% increase in rides at

3. Apps, such as Turo, <https://turo.com/>, are starting to offer similar services.

35% lower prices. Restrictions on driver entry, like those recently adopted in New York City,⁴ reduce welfare, but these losses fall mostly on riders because entry into and exit from ride share employment is flexible.

I begin by exploiting the differential timing of ride share entry across cities in the United States between 2010 and 2016. Ride share entry leads to increases in auto sales, employment, and utilization, despite similar pre-trends in treated and untreated zip codes: auto sales increase by 1.7%; employment, as measured by the number of households filing tax returns, increases by 0.60%; utilization rates increase by roughly 140 miles per day per registered ride share driver. Noting that the average effective full-time wage of an Uber driver is roughly \$23,500 per year, I show that relatively low-income zip codes within treated cities see auto sales increase by 2.6%, as compared to 0.6% for higher-income zip codes. Similarly, the number of tax filings with adjusted gross incomes (AGIs) below \$25,000 increases by 1%, while the number of tax filings above \$25,000 is unchanged. Vehicle utilization increases are confined to low-income zip codes, while high-income zip codes see utilization rates fall. Exploiting ride share vehicle eligibility requirements and vehicle-level micro data, I confirm that the increases in auto registrations and utilization are confined to eligible vehicles in low-income zip codes.

These increases in sales, employment, and utilization are accompanied by spikes in consumer auto loan origination. Treated zip codes see a 1% increase in new loans relative to untreated zip codes with no difference in pre-trends. Like auto sales, originations in low-income zip codes increase by 2% while high-income zip codes see no change. Registration-level micro data reveal that ride share eligible vehicles, in particular, are more likely to be financed than ineligible vehicles. These facts suggest that financing is intimately linked to the real effects of ride share entry through drivers' ability to obtain cars and become ride share drivers.

Given these concurrent increases in financing, I test whether borrowers' inability to borrow short-

4. <https://www.nytimes.com/2018/08/08/nyregion/uber-vote-city-council-cap.html>, accessed November 9, 2018.

circuits their ability to acquire cars and find employment in the gig economy by showing that credit constraints lead to less borrowing and mute the real effects of ride share entry. First, I exploit the exogenous removal of bankruptcy filing history from consumer credit reports. The Fair Credit Reporting Act (FCRA) requires that credit agencies remove bankruptcy information ten years after filing. Borrowers who recently had their bankruptcy filing records expunged face lower borrowing costs than borrowers who will soon have their filing records expunged. I find that ride share entry increases the probability that low financing cost borrowers obtain auto loans, while there is no effect among high financing cost borrowers.

Zip code level evidence confirms this finding. I exploit ex-ante variation in credit access coming from two sources. First, I use zip-level differences in the 2010 share of auto credit from depository institutions, which subsequently saw an increased regulatory burden as regulations arising out of the financial crisis disproportionately affected banks.⁵ Second, I use zip-level differences in 2010 consumer finance delinquency rates, which, consistent with the preceding borrower-level bankruptcy evidence, impaired borrowers' ability to obtain financing years later when ride share entered. I find that zip codes with constrained credit access see essentially no auto loan, sales, or employment growth, as potential drivers cannot get the necessary financing.

My empirical results show how gig economy technologies introduce a fundamental change in the organization of firm production. Unlike a traditional firm that owns capital, hires labor, and sells services, Uber and Lyft coordinate independent drivers who own their own cars, thus transforming consumer durable goods into productive capital. By drawing on a large stock of idle cars already in the economy, Uber and Lyft's technology has the potential to increase the utilization rate, and consequently the effective size of the capital stock. Ride share technology, however, has a critical limitation: it does not match drivers to cars. Rather, ride share relies on the financial system to allocate cars to drivers. In consequence, financial constraints dampen the real effects.

5. See Benmelech et al. (2017) and Buchak et al. (2018a).

The gig economy's unique organizational structure raises three questions. First, the preceding results show that low-income individuals exit the market when finance becomes more expensive or unavailable. How does their exit, together with higher financing costs, impact riders and other drivers who are less dependent on finance? Second, the gig economy currently relies on the financial system to match drivers with cars. This stands in contrast to a traditional taxi company that owns cars and hires drivers, or a hypothetical car-share technology that allows drivers to rent cars on-demand from other households with idle capital. Are there significant differences in terms of ride prices and quantities under these various mechanisms, and do the gig economy's altered firm boundaries increase productivity? Third, many cities have proposed or are currently implementing quantity restrictions on the number of ride share drivers. How large are welfare losses, and how are the losses distributed among drivers and riders when cities adopt these policies?

To assess the impact of these counterfactuals, I build a quantitative structural model that links consumers' vehicle acquisition, utilization, ride share driving, and financing decisions with a simple model of ride share demand. The structural approach complements the reduced-form results by allowing me to quantify equilibrium outcomes arising from several interconnected markets after the introduction of counterfactual organizational structures. Further, the model quantifies outcomes along several dimensions, including ride share wages and policy effects on low- and high-income individuals.

As a starting point, I follow quantitative models of consumer demand in the spirit of Berry et al. (1995) and Nevo (2001), which feature consumers with heterogeneous preferences over price and vehicle characteristics. I extend these models by incorporating financing needs and a ride share labor market decision where consumers with eligible vehicles can decide whether to enter ride share driving. I close the model with a simple model of ride demand. I extend the traditional estimation approach by using rich micro data on vehicle utilization, financing, and ride share driving by matching levels and covariances of these outcomes with consumer demographics.

The model shows that increasing financing costs from the low 2010–2016 levels to historical aver-

ages, such as those during the previous tech boom in the late 1990s, leads to a significantly smaller and less disruptive ride share sector. Total rides decrease by 8%, and prices increase by 20%. This occurs because the low-income individuals comprising the majority of ride share drivers are the most reliant on financing. As financing becomes more expensive, these low-income drivers exit the market, and less finance-reliant drivers enter. These drivers, however, have higher earning potential, and, in consequence, demand higher ride share wages and prices. Costlier financing increases driver net income, even after accounting for higher interest payments, but low-income, finance-dependent drivers do not see these benefits because they have already exited the market. Restricting borrowing entirely leads to drivers with a 50% higher outside earning potential, nearly doubles driver net income, and leads to 18% fewer rides at 25% higher prices.

I next compare the gig economy's mechanism for allocating capital to other structures. First, I consider a *traditional* firm structure—a taxi company—where firms use financial markets to obtain cars, which they allocate to hired drivers. Second, I consider the current *gig economy* structure, where firms match drivers and riders but rely on the financial system to match drivers with cars. Third, I consider a hypothetical *car share* structure similar to currently emerging technologies where households can rent their cars on-demand to ride share drivers. This analysis holds the competitive and regulatory environment fixed and isolates the impact of firm structure and finance. The model shows that the current gig economy structure leads to 22% more rides and 35% lower prices versus the traditional structure, and the third structure—the car share structure—leads to a 20% ride quantity increase and a 30% ride price decrease versus the gig economy structure.

The gains from moving from a traditional economy to a gig economy, and then to a car-share economy, are driven by increases in the supply of ride share capital. First, ride share technologies make it easier for vehicles to serve as both productive capital and consumer durables. Without ride share, households enjoy owning cars for their own consumption but leave them idle most of the time. Technologies that enable drivers to more easily access this idle capital stock increase the effective capital supply. Second, the frictionless rental market technology allows households

with the lowest financing costs to own cars while enabling low-income individuals to drive them. The key lesson is that the gig economy's mechanism of capital allocation has the potential to significantly increase the size and penetration of the gig economy, so long as individuals are able to borrow.

Finally, quantity restrictions on the number of drivers lead to significant welfare losses that fall primarily on riders, rather than drivers. These losses exceed \$300 million per year for a New York City-sized market. Reducing the number of drivers by 25% increases the effective ride price by nearly 50%, including non-price changes, such as increased queuing times. Drivers' effective wages decrease by roughly 5%. Welfare losses fall primarily on consumers because demand is relatively inelastic and supply is relatively elastic: ride share drivers have other similar earning opportunities to which they can easily substitute.

I begin by discussing the data and institutional background. Then, I establish a series of reduced-form facts before moving to the structural model. I conclude by discussing my results and connecting them to the broader literature.

CHAPTER 2

DATA AND INSTITUTIONAL BACKGROUND

2.1 Data

I briefly describe the data used in the analysis. Appendix C.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 that provides 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, where each car is identified by a vehicle identification number (VIN). VINs join to data from the National Highway Traffic Safety Administration and FuelEconomy.gov databases, which provide detailed physical attributes of the car, including make, model, year, horsepower, and fuel efficiency. I use borrower-level detail from TransUnion, which is a 10% sample of all individuals in the United States on which TransUnion has information. It 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, which provides the number of zip-year level tax filings by AGI brackets, and the United States Census and American Community Survey (ACS).

2.2 Institutional Background

Uber began operations in San Francisco in 2010, with Lyft following shortly thereafter. Both services expanded rapidly to other cities. Figure A.1 Panel (a) shows the staggered entry of ride

share across markets by month. The entry rate accelerates quickly in 2014 and 2015. Panel (b) shows the total number of Uber drivers in the United States between 2012 and 2016. By the end 2016, there were nearly 800,000 registered Uber drivers.¹ Panel (c) shows the number of drivers in three representative cities, including San Francisco. In San Francisco, where Uber has had the longest presence, the number of drivers continues to grow rapidly. Panel (d) 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.

The timing of ride share entry is not random. In Appendix Section C.2, I study what predicts ride share entering a city. In short, entry is significantly more likely in large cities with high mobile broadband penetration, consistent with the idea that these services enter areas with large potential markets. Variables such as vehicle ownership rates or access to finance do not predict entry. This is true both on the extensive margin of entry and in terms of how quickly these services enter.

1. 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.

CHAPTER 3

RIDE SHARE'S REAL EFFECTS AND THE ROLE OF FINANCE

This section studies the empirical effects of ride share entry and the role that finance plays in its expansion. I begin with an analysis of how cars were distributed prior to ride share entry. I then study how sales, loans, employment, and utilization respond to entry. I conclude by studying how financial constraints inhibit these effects.

3.1 Who owns cars before ride share?

Other things equal, low-income or marginally employed households benefit more from ride share employment than high-income households because their opportunity cost of time from driving is lower. If low-income households already have cars, they can drive for ride share immediately upon entry. If not, they must acquire cars in order to take up ride share employment. I therefore begin by studying the ex-ante distribution of cars across the income spectrum. Using individual data from the 2010 ACS, I calculate the number of cars per adult household member and run the following household-level regression:

$$CarsPerHH_i = \sum_b \gamma_b I(Income_i \in Bin_b) + \gamma_c + \epsilon_i \quad (3.1)$$

$CarsPerHH_i$ 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. Figure A.2 plots γ_b versus the income bin Bin_b . Panel (a) shows raw means across incomes while Panel (b) includes the CBSA fixed effect.

These figures show that vehicle ownership rates are strongly positively 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. This is true across both specifications. In Appendix Section C.3, I run a similar regression at the zip code level and include a richer set of covariates. I confirm that the strong relationship between income and ex-ante vehicle ownership persists after controlling for commute time, mobile broadband access, and other demographic variables.

3.2 What happens when ride share enters?

Ride share driving offers relatively low wages, and so low-income households are the most natural drivers. The preceding analysis showed, however, that low-income households were less likely to own cars before ride share entered. Therefore, many low-income households will need to acquire them in order to become ride share drivers. In this section, I study how ride share entry impacts auto sales and loans, how these effects differ by high- and low-wage households, and whether these effects carry through to employment growth and utilization increases as individuals become professional ride share drivers.

3.2.1 Auto sales and loans

My empirical approach exploits ride share's staggered entry. Treatment occurs at the city level. I first run a difference in difference analysis using the nationwide RL Polk and Equifax datasets comparing auto sales and loans in zip codes that ride share has entered to zip codes that ride share has not yet entered. I then exploit within-city zip code variation in median incomes to test whether the impact of treatment is concentrated in relatively low-income zip codes. The main specifications

are as follows:

$$\log Auto_{zt} = \sum_{\tau=-4}^4 \beta_{\tau} ET_{zt}^{\tau} + \gamma_t + \gamma_z + \epsilon_{zt} \quad (3.2)$$

$$\log Auto_{zt} = \beta Post_{zt} + \gamma_t + \gamma_z + \epsilon_{zt} \quad (3.3)$$

$$\log Auto_{zt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Low Income_z + \gamma_z + \gamma_{Income,t} + \epsilon_{zt} \quad (3.4)$$

Equation (3.2) is the event study; Equations (3.3) and (3.4) are the difference in difference and within-city analyses, respectively. $Auto_{zt}$ is new vehicle sales from RL Polk, or new auto loan originations from Equifax in zip z at time t . ET_{zt}^{τ} is an indicator for whether ride share entered zip z τ periods before time t . $Post_{zt}$ is an indicator for whether ride share has entered zip z on or before time t . $Low Income_z$ is an indicator for whether the zip code's median income is in the bottom 50% for zip codes in the CBSA.¹ 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. These regressions include only zip codes that receive treatment at some point; those zip codes that are untreated over the entire sample are excluded.

Figure A.3 Panels (a) and (b) show the event study. Panel (a) shows new auto sales and Panel (b) shows new auto loans. This figure plots the coefficients β_{τ} on the y-axis with event time τ on the x-axis, and captures how auto sales or loans deviate from the zip-level mean around ride share entry, adjusted for nationwide time trends.² Panel (a) shows that prior to ride share entry, there are no pre-trends in new auto sales. Once ride share enters, auto sales begin to increase significantly. Panel (b) shows that these increases in auto sales are accompanied by increases in auto loans, which follow a nearly identical pattern. Note that the pattern of growth closely follows the dynamics of driver entry shown in Figure A.1 Panel (d): growth is initially slow and then begins to increase

1. The prediction is that ride share drivers should come from households that have low outside incomes relative to others in the city, rather than low incomes in absolute terms.

2. 95% confidence intervals are also shown. Standard error bands are calculated by bootstrapping the regression by randomly selecting the sample of treated CBSAs.

rapidly in the quarters following ride share entry.

I next split the sample into high- and low-income zip codes and rerun the event study. These results are shown in Figure A.4 Panels (a) and (b) for sales and (c) and (d) for loans. Panels (a) and (c) show the results for low-income zip codes. Panels (b) and (d) shows the results for high-income zip codes. Auto sales in low- and high-income zip codes follow similar pre-trends, but following entry, sales increase significantly in low-income zip codes while remaining essentially unchanged in high-income zip codes. As before, the results for auto loans follow a nearly identical pattern. These results confirm that auto sales and loans increases occur precisely among the demographics most likely to take up ride share driving. Interestingly, high-income zip codes do not reduce their purchases, suggesting that these households value the convenience of car ownership even after ride share services become available.

The difference in difference regressions confirm these results. Table B.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.

Robustness checks: In addition to the event studies showing that there are no pre-trends in auto sales prior to ride share entry, I perform three additional robustness checks. First, I exploit ride share vehicle eligibility requirements and detailed vehicle-level registration data. In particular, 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. 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. Appendix Section C.4 provides considerable detail regarding this analysis. This analysis both provides clear evidence that the growth in auto sales and loans is driven by entry into ride share driving and highlights how ride share entry transforms consumer durables into productive capital. Ride share-productive capital (eligible cars) move increasingly into the hands of ride share-productive (low-income) households.

As a second robustness check, to insure that splitting the sample by zip-level income is picking up variation that is relevant to ride share *driving* in particular, I repeat the preceding analysis but split zip codes on 2010 transportation worker share rather than wage. The results, shown in Column (3) of Table B.1 Panels A and B, show that auto sales and loan increases are concentrated in zip codes with a high share of transportation workers. This confirms that the types of households that increase purchases after ride share entry are those that previously worked as drivers in another industry. Finally, I run a placebo test of Regressions (3.3) and (3.4) where the timing of ride share entry is held fixed but assigned randomly across locations. This test, shown in Appendix Table C.7 Panels A and B, finds no significant effect, confirming that the preceding findings are related to ride share entry and not unrelated trends in the data. In Appendix Section C.5, I show that loan performance is unaffected by ride share entry.

3.2.2 *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 out by adjusted gross income (AGI). I run the following event study and

difference in difference regression at the zip-year-AGI bracket level:

$$\log Filings_{bzt} = \sum_{\tau=-3}^3 \beta_{\tau} ET_{zt}^{\tau} + FE + \epsilon_{bzt} \quad (3.5)$$

$$\log Filings_{bzt} = \beta_1 Post_{zt} + \beta_2 Post_{zt} \times I(AGI_b < 25k) + FE + \epsilon_{bzt} \quad (3.6)$$

$Filings_{bzt}$ is the number of tax filings in AGI bracket b in zip z at year t . ET_{zt}^{τ} is an indicator for whether ride share entered zip z τ periods before time t . $Post_{zt}$ is an indicator for whether ride share has entered zip code z on or before year t . FE is a set of fixed effects. The event study includes year and zip fixed effects, while the difference in difference regression includes zip \times AGI bracket and time \times AGI bracket fixed effects. These regressions include only zip codes that receive treatment at some point in the sample period.

The results show that ride share entry leads to greater employment as measured by tax filings. Figure A.5 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 by roughly 0.5%. Panels (b) and (c) split the sample into filings below \$25,000 and filings above \$25,000, respectively. The yearly income of a full-time ride share driver is approximately \$23,500, so households entering the labor market as drivers will for the most part file in the \leq \$25,000 AGI bracket. These event studies confirm that tax filings increase by roughly 1% among low-AGI filers with no change among high-AGI filers. These quantities match closely with the percentage of ride share drivers in cities in Figure A.1 Panel (d).

The difference in difference regression confirms these results. Table B.2 Column (1) shows an increase in tax filings across all AGI brackets of roughly 0.7% following ride share entry. Column (2), which includes the AGI interaction, shows that the increase in filings is concentrated entirely in filings in the \leq \$25,000 AGI bracket. 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. I run a placebo test where

the location of ride share entry is randomized across cities. These results are shown in Table C.9 and show no effect.

3.2.3 Vehicle utilization

The preceding analysis shows that ride share entry increases auto sales and loans among low-income households and employment among low-income tax filers. This suggests that households acquire cars and enter gig economy employment. I next test whether these cars see higher utilization rates after entry. The South Carolina DMV data report vehicle odometer readings when the car changes ownership. I calculate yearly utilization rates for completed ownership spells as the difference in miles at the start and end of the spell divided by the spell length. I then study how utilization rates change following ride share entry and whether there is a differential effect for eligible and ineligible vehicles across high- and low-income zip codes. I run the following regression:

$$\begin{aligned} MilesPerYear_{vzt} = & \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Eligible_v + \beta_3 Post_{zt} \times Low\ income_z \\ & + \beta_4 Post_{zt} \times Eligible_v \times Low\ income_z + FE + \epsilon_{vzt} \end{aligned} \quad (3.7)$$

The term FE collects various fixed effects, the details of which are described below. Table B.3 shows the results. Column (1), which shows the overall effect of ride share entry, contains zip times eligibility and quarter times eligibility fixed effects. This column shows a positive but statistically insignificant effect. Adjusting for the number of ride share drivers, this number implies that the typical ride share driver drives roughly 140 miles per day.³

Column (2) adds the $Post \times Low\ income$ interaction term with the same sets of fixed effects as in Column (1). This regression finds a larger but still insignificant effect in low income zip codes of roughly 400 miles per year. Column (3) includes a $Post \times Eligible$ interaction and

3. 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.

finds a large and statistically significant difference of roughly 1,350 more miles per year in eligible vehicles versus ineligible vehicles following ride share entry. This is offset by large and statistically significant decrease in ineligible vehicle utilization of roughly 1,000 miles per year. Column (4) repeats the analysis in Column (3) with the inclusion of $\text{zip} \times \text{quarter}$ fixed effects. These fixed effects absorb any systematic changes in overall mileage driven in a particular zip code, including, for example, additional working or socializing opportunities. The differential impact between eligible and ineligible vehicles is largely unchanged.

Columns (5) and (6) add the triple interaction of $\text{Post} \times \text{Low income} \times \text{Eligible}$. Column (5) has $\text{zip} \times \text{Eligible}$ and $\text{quarter} \times \text{Eligible}$ fixed effects. Column (6) adds $\text{zip} \times \text{quarter}$ fixed effects. Both columns show a large and statistically significant differential effect for eligible vehicles in low-income zip codes of more than 2,000 additional miles per year compared to ineligible vehicles. Figure A.6 shows the corresponding event study. This figure shows the within zip difference in utilization between eligible and ineligible vehicles versus event time. The utilization rate increases at the time that ride share enters, while there are no pre-trends in utilization rates prior to entry. Table C.10 shows the results of a placebo test where treatment is randomized among CBSAs, and finds no effect.

I conclude this section by contrasting the dynamics of vehicle utilization to the dynamics of auto sales, loans, and employment. Figures A.3–A.4 show gradual increases in auto sales and tax filings following ride share entry. These dynamics roughly match the rate of driver entry shown in Figure A.1 Panel (d). The event study for utilization shown in Figure A.6, on the other hand, shows that utilization increases immediately. This shows that those individuals living in low income zip codes who already possess ride-share eligible cars begin driving for ride share immediately. This leads to a utilization increase on impact among eligible cars in low-income zip codes. Sales, loans, registrations, and tax filings, in contrast, increase more slowly as households adjust their capital purchases. These findings show how the productivity of the capital stock, measured as utilization rates, increases through factor reallocation. Here, ride-share eligible vehicles are more productive

in the hands of low-income households following ride-share entry, who acquire them after entry and use them more.

3.3 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. Ride share entry creates employment opportunities for low-income households if they possess the necessary capital. Low-income households that already possess the capital increase utilization, while other households borrow and acquire the necessary capital, use it for ride share driving, and start earning taxable income. In this section I study the consequences for the gig economy if this process is short-circuited by households being unable to borrow.

I approach this question in two ways. First, I use borrower level data from TransUnion to generate exogenous variation in borrower credit access. A federal statute requires credit agencies to remove Chapter 7 bankruptcy filings from credit reports no more than ten years after filing. This generates variation in borrowing costs between otherwise similar borrowers who have or have not yet had this bankruptcy information removed. I show that borrowers facing greater costs borrow less in response to ride share entry. Second, I use zip-level variation in credit access to study how the response of auto lending, sales, and employment to ride share entry varies with credit access. 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. I show that these measures predict less auto loan growth and lead to smaller real effects of ride share entry.

3.3.1 Borrower-level evidence from bankruptcy flag removal

I begin by exploiting the exogenous removal of bankruptcy filing information from consumer credit records. In particular, following Musto (2004) and more recently Herkenhoff et al. (2016) and Dobbie et al. (2016), I use the fact that the Fair Credit Reporting Act (FCRA) requires credit bureaus to remove Chapter 7 bankruptcy filing reports no more than ten years after filing. Previous literature has documented that this exogenous removal leads to a large and statistically significant increase in credit scores and a decrease in borrowing costs.

The data in this section come from TransUnion borrower-level credit reports. I begin by showing that the bankruptcy flag's removal is relevant for the level of auto lending, independent of ride share entry. Figure A.7 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 regressions to isolate the impact of filing removal:

$$Origination_{izt} = \sum_{\tau=7.5}^{12.5} \beta_{\tau} I(YearsSinceFiling_{izt} = \tau) + \gamma_{zt} + \epsilon_{izt} \quad (3.8)$$

$$Origination_{izt} = \beta I(YearsSinceFiling \geq 10) + \gamma_{zt} + \epsilon_{izt} \quad (3.9)$$

Regression (3.8) compares auto loan origination for consumers who have filed for bankruptcy at different times in the relatively distant past in the same zip code at the same time. Regression (3.9) is a similar approach but simply compares borrows above and below the 10 year cutoff. Figure A.7 Panel (b) shows the coefficients β_{τ} from (3.8), and finds a clear discontinuity occurring at the point where the flag is required to be removed, with no systematic differences occurring before or after the cutoff. Table B.4 Panel A shows the results for regression (3.9). 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.

Next, to study how credit access translates 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. Figure A.8 illustrates the timing of this experiment. I run the following regression:

$$Origination_{i,z,t} = \beta_1 Post_{z,t} + \beta_2 Post_{z,t} \times Constrained_i + \gamma_{gt} + \gamma_{gz} + \gamma_{zt} + \epsilon_{i,z,t} \quad (3.10)$$

The regression is at the borrower-quarter level. $Origination_{i,z,t}$ is an indicator for whether the borrower obtains an auto loan in the quarter. $Post_{z,t}$ 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_{gz} 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 B.4 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.⁴ Column (2) includes the interaction of treatment with whether the bankruptcy flag is

4. 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

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, limiting the sample of borrowers to 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 A.7 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 significant effects leading up to entry. This analysis shows that borrowers facing greater credit costs increase borrowing less in response to ride share entry.

3.3.2 *Zip-level evidence*

I next use 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. An auto loan generally comes from a bank or a captive lender like GM financial. Following the financial crisis, new banking regulations such as increased capital requirements and stricter super-

not declared bankruptcy.

vision reduced banks' ability to lend.⁵ Those zip codes where banks dominated auto lending in 2010 saw their lenders hampered by regulations and face a more constrained supply of auto credit relative.⁶

The empirical strategy mirrors regressions (3.2) and (3.4). Rather than splitting by zip-level income, however, I split by the ex-ante bank share and the ex-ante consumer finance delinquency rate. I begin by studying loan originations. Figure A.9 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.

Regressions following equation (3.3) 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. The left-hand side variable is log auto loan originations. Table B.5 Panel A shows the results. 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 A.9, 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.

5. 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).

6. 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.

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 A.10 for auto sales and A.11 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)).

Finally, Table B.5 Panel B shows the corresponding regressions. 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. I conclude with a placebo test, shown in Appendix Table C.11, which randomizes the location of ride share treatment and shows no effect.

These results establish that credit constraints reduce the impact of ride share entry on new loan origination. Fewer originations on entry leads to smaller real effects as the borrowers who would like to acquire cars to join ride share employment are unable to do so. These results establish a causal 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.

CHAPTER 4

A STRUCTURAL MODEL OF RIDE SHARE CAPITAL

The reduced form analysis showed that ride share relies on the financial system to help low-income individuals obtain financing to purchase cars. Where financial constraints prevent borrowing, the real effects of gig economy entry are smaller, and the potential productivity gains from Uber and Lyft's technology are limited. This observation raises three questions. First, given that Uber and Lyft rely on the financial system to allocate capital to low-income drivers, what are the consequences of the broader credit environment in which these apps arose for ride share growth and individuals' entry into ride share employment across the income spectrum. In particular, Uber and Lyft arose in a time of exceptionally low interest rates and access to auto credit. Had they arisen in a more typical credit environment, would these technologies have been less disruptive, would low-income individuals have been excluded from the market, and how would driver wages differ?

Second, how large would the gig economy be if a mechanism other than consumer finance allocated cars to drivers? Consider two alternatives. First, a *traditional* structure, where firms finance the acquisition of cars and hire drivers. Second, a *car share* structure that allows drivers to frictionlessly rent cars from other households on an hourly basis. I study whether the gig economy's disaggregated structure itself contributed to its growth and whether further disaggregation through the a frictionless rental market would grow it further.

Third, many cities are implementing or are considering implementing quantity restrictions on the number of drivers. How large are welfare losses from these restrictions, and are they primarily borne by drivers or riders? To assess the impact of these counterfactuals, I build and estimate a quantitative structural model that links consumers' vehicle acquisition, utilization, ride share driving, and financing decisions with a simple model of ride share demand.

4.1 Model description

Consumer $i \in \{1, \dots, I\}$ living in market $c \in \{1, \dots, C\}$ makes three decisions: (1) which (if any) vehicle $j \in \{0, 1, \dots, J\}$ 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 , with the precise mapping between D_i and θ_i to be estimated from the data.

A market c is at the zip-quarter level and is summarized by a distribution of consumer demographics, $F_c(D_i)$, from which individual demographics D_i are drawn, and ϕ_c , an indicator with $\phi_c = 1$ when ride share is present in the market and $\phi_c = 0$ otherwise. Each vehicle $j \in \{0, 1, \dots, J\}$, with $j = 0$ being the outside option of not owning a vehicle, has physical attributes x_j . These attributes include, for example, price, body type, gas mileage, and horsepower. Additionally, the vehicle has ride-share eligibility status e_j , with $e_j = 1$ if the vehicle is eligible to be driven for ride share and $e_j = 0$ otherwise. Finally, the vehicle's price, p_j , can be paid upfront provided the consumer has sufficient liquidity, or can be financed at a cost regardless of the consumer's cash liquidity.

To close the model, I include a simple model of ride share demand using demand elasticity estimated in Cohen et al. (2016). The consumer's vehicle choice described above generates a supply of ride share services for a given ride share price, and I calculate the equilibrium by finding the ride share price that equalizes supply and demand for ride share services. I estimate the parameters governing vehicle choice, financing choice, and driving choice separately from the demand side of the model, meaning that these estimates are independent of my demand-side assumptions.

4.1.1 Individual consumer's problem

The consumer obtains utility from three sources: (1) utility from owning the car for the purposes of her own use and consumption, (2) disutility from purchasing and financing the vehicle, and (3) utility from driving for ride share. Utility from each source drives a separate decision regarding (1) which vehicle to own, (2) whether to purchase outright or finance the vehicle, and (3) whether to drive for ride share. As in Berry et al. (1995), I directly model indirect utility. I first describe these sources of utility separately before stating the consumer's decision problem and characterizing her optimal choices.

Utility from own consumption

Owning vehicle j provides consumption-like utility arising from (1) the vehicle's characteristics and (2) the consumer's utilization of the vehicle for her own driving purposes. In particular, own consumption indirect utility takes the following linear form:

$$u_c(\theta_i, x'_j, \phi_c, \xi_{jc}, \epsilon_{ijc}^c) = \underbrace{\beta_i^0}_{\text{outside option}} + \underbrace{x'_j \beta_i + \xi_{jc}}_{\text{vehicle characteristics}} + \underbrace{\beta_i^m m_i}_{\text{utilization}} + \underbrace{\epsilon_{ijc}^c}_{\text{preference shock}} \quad (4.1)$$

β_i^0 captures outside option value of not buying a car, which is common across all vehicles. β_i^0 arises in the indirect utility function from the consumer's budget constraint and is allowed to vary before and after entry. The vehicle characteristics term captures utility the consumer gets from owning a car with certain characteristics. x'_j contains car characteristics observed to both consumer and econometrician such as body type, fuel efficiency, and horsepower. ξ_{jc} summarizes car characteristics that are observed by the consumer but not by the econometrician, such as the vehicle's brand. β_i is a vector of consumer preferences, elements of consumer characteristics θ_i , over the observed car characteristics. These summarize, for example, how desirable additional horsepower is in the vehicle. These preferences vary systematically across demographics, meaning, for example, that

higher wage consumers may prefer cars with larger engines.

The utilization term captures utility the consumer gets from using the car. m_i is the consumer's desired miles driven. m_i , like other consumer characteristics, varies systematically with consumer demographics, and with whether ride share has entered the market. This allows the model to capture the fact that own utilization may change when ride share enters a market. β_i^m is the consumer's utility per mile driven. Finally, ϵ_{ijc}^c is an idiosyncratic shock to consumer i 's preference for vehicle j . This captures the fact that otherwise identical consumers may choose different cars for reasons that are not explicitly modeled. For example, consumer i may choose car j because she lives close to the dealership.

Utility from financing

Consumer i can buy a car in two ways: Out of pocket or with a loan. Her decision depends on (1) the vehicle's price, (2) her available liquidity, and (3) her financing costs. Vehicle j costs p_j . Regardless of her financing decision, the consumer receives disutility $\alpha_i p_j$ from acquiring the car. $\alpha_i > 0$ is her price sensitivity, and other things equal she prefers a cheaper car. In addition to disutility arising directly from the vehicle price, consumer i decides whether or not to obtain financing. The indirect disutility of obtaining financing has the following functional form:

$$f(\theta_i, p_j, \epsilon_{ijc}^f) = f_i^0 + \epsilon_{ijc}^f \quad (4.2)$$

f_i^0 captures fixed costs in obtaining financing, for example, search costs or origination fees.¹ f_i^0 varies with consumer demographics, so this specification accommodates, for example, high-

1. In unreported results, I estimate a specification where financing costs depend linearly on the amount the borrower must borrow, with

$$f(\theta_i, p_j, \epsilon_{ijc}^f) \equiv f_i^0 + f_i^1(p_j - l_i) + \epsilon_{ijc}^f$$

Here, f_i^1 captures costs that are increasing in the amount she must borrow relative to what she can put down in cash. For example, a vehicle loan with less money down might bear a higher interest rate. This more flexible specification does not improve the model fit or meaningfully change the results, so I report the results using the simpler specification.

income borrowers having access to lower-cost credit than low-income borrowers. ϵ_{ijc}^f is a consumer-specific shock to financing utility, capturing shocks that differ across demographically similar consumers. For example, consumers may have idiosyncratic financing or consumption needs, differing interest rate elasticities, or financial sophistication, that make vehicle financing more appealing.

Financing is required if the consumer's liquidity l_i is insufficient to cover the vehicle's price outright, i.e., if $l_i < p_j$. In that case, the financially constrained consumer necessarily suffers the disutility of obtaining financing if she selects the vehicle. If the consumer has sufficient liquidity to pay without financing, she can choose whether to obtain financing. The following equation summarizes the consumer's disutility from obtaining a car, depending on her constraints and whether she obtains financing:

$$u_f(\theta_i, x'_j, \epsilon_{ijc}^f) = -\alpha_i p_j - \begin{cases} f(\theta_i, p_j, \epsilon_{ijc}^f) & l_i < p_j \text{ (constrained)} \\ \begin{cases} f(\theta_i, p_j, \epsilon_{ijc}^f) & \text{financed} \\ 0 & \text{not financed} \end{cases} & l_i \geq p_j \text{ (unconstrained)} \end{cases} \quad (4.3)$$

Utility from ride-share driving

Consumer i can use her car to 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:

$$u_{rs}(p, \theta_i, x'_j, \phi_c, \epsilon_{ijc}^f) = \begin{cases} \begin{cases} \alpha_i (p\zeta - b - w_i) + \gamma_i + \epsilon_{ijc}^d & \text{drive} \\ 0 & \text{do not drive} \end{cases} & e_j, \phi_c = 1 \\ 0 & \text{otherwise} \end{cases} \quad (4.4)$$

p is the price that ride share riders pay for rides. ζ is the portion of the price that drivers keep after the app's commission,² b represents expenses such as maintenance and foregone full-time employee benefits. p is determined in equilibrium. w_i is the individual's outside wage. $(p\zeta - b - w_i)$ captures the total economic 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. Scaling this term by price elasticity α_i puts it in utility terms. γ_i captures the non-monetary net benefits of driving, such as the utility from having flexible hours, and her utility from leisure.³ ϵ_{ijc}^d is an idiosyncratic shock that captures consumers with observably similar economic benefits differing in their driving utility. For example, a parent may dislike driving at peak hours due to child care obligations.

Combined utility

The consumer's total indirect utility is the sum of her utility from own consumption, financing, and ride share driving:

$$u(p, \theta_i, x'_j, \phi_c, \xi_{jc}; \epsilon_{ijc}^c, \epsilon_{ijc}^{rs}, \epsilon_{ijc}^d) = \underbrace{u_c(\theta_i, x'_j, \phi_c, \xi_{jc}, \epsilon_{ijc}^c)}_{\text{own consumption}} + \underbrace{u_f(\theta_i, x'_j, \epsilon_{ijc}^c)}_{\text{financing}} + \underbrace{u_{rs}(p, \theta_i, x'_j, \phi_c, \epsilon_{ijc}^d)}_{\text{ride share driving}} \quad (4.5)$$

2. For example, if $\zeta = 0.75$ then the app takes 25% of the ride fee. Uber also charges a ride booking fee of approximately \$1.50 per ride. This money is charged to the rider when requesting a ride.

3. Note that γ_i captures the intensive margin of ride share driving. Observe that a consumer allocating her time between leisure, ride share work, and non-ride share work would allocate all of her working time towards the higher paying job and then allocate time between work and leisure. A consumer who values leisure time relatively more would obtain less value from driving for ride share, which is captured in γ_i .

Optimal consumer choice

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.⁴ I characterize the solution to the financing and driving discrete choice problems before solving the overall vehicle choice problem.

Financing choice

Recall from (4.3) 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:

$$u_f^*(\theta_i, x'_j, \epsilon_{ijc}^f) = -\alpha_i p_j + \begin{cases} -f(\theta_i, p_j, \epsilon_{ijc}^f) & l_i < p_j \\ \max_{f, \sim f} \{-f(\theta_i, p_j, \epsilon_{ijc}^f), 0\} & l_i \geq p_j \end{cases} \quad (4.6)$$

4. 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 . I justify the financing assumption by the fact that the consumer receives her specific financing and loan terms after she has decided which vehicle she likes based on its characteristics. I justify the ride share assumption by the fact that the driver does not know how she will enjoy the idiosyncratic, non-monetary aspects of ride share driving, such as flexibility in working hours or socialization with passengers, until she actually tries ride share driving.

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:

$$p_f(\theta_i, x'_j) = \begin{cases} 1 & l_i < p_j \\ \frac{\exp(-f(\theta_i, p_j))}{\exp(-f(\theta_i, p_j)) + 1} & l_i \geq p_j \end{cases} \quad (4.7)$$

$$u_f(\theta_i, x'_j) = -\alpha_i p_j + \kappa + \begin{cases} -f(\theta_i, p_j) & l_i < p_j \\ \log [\exp(-f(\theta_i, p_j)) + 1] & l_i \geq p_j \end{cases} \quad (4.8)$$

where κ is the Euler-Mascheroni constant. $p_f(\theta_i, x'_j)$ denotes the probability that consumer i finances the vehicle conditional on obtaining it; $u_f(\theta_i, x'_j)$ denotes the consumer's ex-ante expected utility from purchasing and financing the vehicle at the time she makes the vehicle acquisition decision.

Ride share choice

Recall from (4.4) 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:

$$u_{rs}^*(p, \theta_i, x'_j, \phi_c, \epsilon_{ijc}^{rs}) = \begin{cases} \max_{d, \sim d} \left\{ \alpha_i (p\zeta - b - w_i) + \gamma_i + \epsilon_{ijc}^d, 0 \right\} & e_j, \phi_c = 1 \\ 0 & \text{otherwise} \end{cases} \quad (4.9)$$

With the assumption that ϵ_{ijc}^{rs} follows a type-one extreme value distribution, the probability that she drives for ride share, and the expected utility from the ride share decision are:

$$p_{rs}(p, \theta_i, x'_j, \phi_c) = \begin{cases} \frac{\alpha_i(p\zeta - b - w_i) + \gamma_i}{\exp(\alpha_i(p\zeta - b - w_i) + \gamma_i) + 1} & e_j, \phi_c = 1 \\ 0 & \text{otherwise} \end{cases} \quad (4.10)$$

$$u_{rs}(p, \theta_i, x'_j, \phi_c) = \begin{cases} \log [\exp(\alpha_i(p\zeta - b - w_i) + \gamma_i) + 1] + \kappa & e_j, \phi_c = 1 \\ 0 & \text{otherwise} \end{cases} \quad (4.11)$$

$p_{rs}(p, \theta_i, x'_j, \phi_c)$ denotes the probability that consumer i drives ride share conditional on obtaining the vehicle; $u_{rs}(p, \theta_i, x'_j, \phi_c)$ denotes the consumer's ex-ante expected utility from driving ride share at the time she makes the vehicle acquisition decision.

Vehicle choice

I come now to the vehicle choice. Her indirect utility from choosing vehicle j is:

$$u(p, \theta_i, x'_j, \phi_c, \xi_{jc}, \epsilon_{ijc}^c) = u_c(\theta_i, x'_j, \phi_c, \xi_{jc}) + u_f(\theta_i, x'_j) + u_{rs}(p, \theta_i, x'_j, \phi_c) + \epsilon_{ijc}^c \quad (4.12)$$

Her vehicle choice solves the following discrete choice problem:

$$\max_{j \in \{0, 1, \dots, J\}} \{u(p, \theta_i, x'_j, \phi_c, \xi_{jc}, \epsilon_{ijc}^c)\} \quad (4.13)$$

Given the logit structure of ϵ_{ijc}^c , the ex-ante probability of consumer i choosing car j is

$$p_{ijc} \equiv p(p, \theta_i, x'_j, \phi_c, \xi_{jc}; \{(x'_k, \xi_{kc})\}) = \frac{\exp(u(p, \theta_i, x'_j, \phi_c, \xi_{jc}))}{1 + \sum_{k=1}^J \exp(u(p, \theta_i, x'_k, \phi_c, \xi_{kc}))} \quad (4.14)$$

4.1.2 Consumer preferences and aggregation

The preceding analysis characterized the decision problem for an individual consumer endowed with characteristics θ_i . This section adds structure to consumer preferences and aggregates individual decisions to market-level predictions.

Characteristics and demographics

I assume that consumer characteristics θ_i have the following distribution:

$$\theta_i = \bar{\theta} + (D_i - \bar{D})'\Pi + \Sigma\epsilon_i^\theta \quad (4.15)$$

$\bar{\theta}$ 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 preferences. 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. Additionally, I allow the characteristic means $\bar{\theta}$ to vary by whether ride share has entered the market.

A key parameter is Π , which describes how demographic translate into individual characteristics and consequently consumer behavior. For example, one element of Π governs how a consumer's outside wage impacts l_i , the consumer's available liquidity. A negative relationship between outside wage and liquidity implies that low outside option wages consumers who value ride share driving will require financing in order to obtain vehicles and drive.

Individual demographics are themselves drawn from a market-level distribution, $F_c(D_i)$:

$$D_i \sim F_c(D_i) \quad (4.16)$$

Together, this gives the key set of structural parameters to be estimated as $\Theta = (\bar{\theta}, \Pi, \Sigma)$, together with the distribution $F_c(D_i)$ which I estimate separately from the main BLP procedure.

Market shares, utilization, driving, and financing

Equations (4.14), (4.10), and (4.7) together characterize the decisions of an individual consumer conditional on her preferences θ_i . Let $A_j(p, x', \phi_c, \xi_c; \Theta)$ denote the set of consumer characteristics such that borrowers with these characteristics in market c choose car j :

$$A_j(p, x', \phi_c, \xi_c; \Theta) = \{(\theta_i, \epsilon_{ijc}^c) | u(p, \theta_i, x'_j, \phi_c, \xi_{jc}, \epsilon_{ijc}^c) > u(p, \theta_i, x'_k, \phi_c, \xi_{kc}, \epsilon_{ikc}^c), \forall k \neq j\} \quad (4.17)$$

Then, the market share of vehicle j in market c given parameters Θ is

$$s_{jc}(p, x', \phi_c, \xi_c; \Theta) = \int_{A_j} p(p, \theta_i, x'_j, \phi_c, \xi_{jc}; \{(x'_k, \xi_{kc})\}) dF(\theta_i; \Theta) \quad (4.18)$$

where $p(p, \theta_i, x'_j, \phi_c, \xi_{jc}; \{(x'_k, \xi_{kc})\})$ is given by (4.14). Following the vehicle choice j , the conditional share of consumers taking up ride share and vehicle financing is:

$$s_{jc}^{rs}(p, x', \phi_c, \xi_c | j; \Theta) = \int_{A_j} p^{rs}(p, \theta_i, x'_j, \phi_c) dF(\theta_i | A_j; \Theta) \quad (4.19)$$

$$s_{jc}^f(x', \phi_c, \xi_c | j; \Theta) = \int_{A_j} p^f(\theta_i, x'_j) dF(\theta_i | A_j; \Theta) \quad (4.20)$$

4.1.3 Demand for ride share services

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

$$\log q = \delta_0 - \delta_1 \log p \quad (4.21)$$

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 in the ride share market is defined as the price p such that the total quantity of ride share services supplied, $s_{jc}^{rs}(p, x', \phi_c, \xi_{\cdot c}|j; \Theta)$ equals the total quantity of ride share services demanded. The main portion of the model generates a supply curve for ride share services, and the counterfactuals I study shift supply in various ways. The demand portion of the model allows me to disentangle how these supply shifts translate to changes in quantities and prices. Appendix Section C.6 studies the sensitivity of my results to this demand specification by varying price elasticity and the ride share commission rate.

4.1.4 Estimation

Estimation strategy

The estimation broadly follows the procedure in Berry et al. (1995) with several additions that utilize model predictions and micro-data regarding utilization, ride share take-up, and financing. In this procedure, given parameters Θ and a vector of vehicle-market unobservables $\xi_{\cdot c}$, I calculate model-predicted market shares by evaluating (4.18) numerically. By inverting this equation, one can obtain a vector $\xi_{\cdot c}^*(\Theta)$ such that given parameters Θ , model-predicted market shares exactly match those in the data. A different Θ' yields a different vector $\xi_{\cdot c}^*(\Theta')$. The estimation procedure iterates over Θ so that a vector of moment conditions involving $\xi_{\cdot c}^*(\Theta')$ are as close to zero as possible. These moments include minimizing the squared sum of $\xi_{\cdot c}^*(\Theta')$, plus standard orthogonality moments, such as $\xi_{\cdot c}^*(\Theta')$ and all vehicle observables having zero covariance.

I extend the standard estimation procedure by bringing additional model predictions to the micro-data. The key parameters are the non-monetary value of driving, the cost of finance, and how the cost of finance varies with outside option income. The average non-monetary value of driving,

$\bar{\gamma}$, governs the overall number of drivers that choose to enter ride share. Increasing $\bar{\gamma}$ leads to more predicted drivers. The estimation therefore identifies $\bar{\gamma}$ by matching the number of drivers observed in the data to the number generated by the model.

The average financing cost, \bar{f}^0 governs the overall level of financing. A higher \bar{f}^0 leads to fewer households obtaining financing. The estimation identifies \bar{f}^0 by matching the predicted level of auto financing to that observed in the data. Financing costs also vary by the individual's outside option income. This relationship is determined by the parameter $\Pi_{f^0,w}$. A negative $\Pi_{f^0,w}$ means that lower income borrowers face higher financing costs. A more negative $\Pi_{f^0,w}$ leads to a steeper predicted relationship between income and financing. I identify $\Pi_{f^0,w}$ by matching this predicted relationship to the relationship between wage and financing.

Finally, several key parameters, such as the hourly ride share price p , the ride share commission rate ζ , and the costs of ride share driving b , I take from Mishel (2018).⁵ For the demand model, I use $\delta_1 = 0.57$ from Cohen et al. (2016).

Estimated parameters

Table B.6 shows the model parameters, with Panel A showing the estimated parameters 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 -1.5 with a standard deviation of 0.075. This corresponds to an average price sensitivity of 0.22, or a mean price elasticity of roughly 4.5, in line with estimates from Berry et al. (1995).⁶ Next, focusing only on the post-entry parameters, the convenience of owning a car is valued at roughly \$7.2 thousand dollars. This indicates that independent of usage, consumers obtain utility

5. 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.

6. Price elasticity in this model is approximately $\frac{\alpha}{1-s} \times p$; the average market share is approximately 0 and the average car price is approximately 20 thousand.

simply through the convenience of owning the car.

Consumers value ride share eligibility at \$5.7 thousand dollars. They additionally value greater horsepower and more fuel efficient vehicles. A one standard deviation increase in horsepower is valued at \$11 thousand, and a one standard deviation increase in fuel efficiency is valued at roughly \$3.2 thousand. This is consistent with households valuing higher quality cars. Financing costs are significant, estimated to be between \$1 and \$2 thousand dollars per year. These costs increase significantly for low-income households, with a 1% decrease in outside option income corresponding to a 1.5% increase in financing costs. Similarly, high-income households have significantly more liquidity, meaning that they require less financing. A household with 1% more income has on average 1.5% greater liquidity with which to purchase a car.

I evaluate the critical relationship between income and financing behavior in Figure A.12. This figure plots the actual and estimated relationship between financing and wage in four subsamples of the data. In each Panel, the dashed line is the model prediction, and the dotted line is the data. Panel (a) shows the results among eligible vehicles, showing a close match between the model and the data, with lower wage households acquiring more financing. The remaining panels show a similar fit. The two most critical subsamples, financing among eligible vehicles, and financing post ride-share entry, exhibit a particularly close match between model and data, finding a negative relationship between wage and financing.

4.2 Counterfactuals

I first study how financing costs affect the size, prices, distribution of capital, and employment in the gig economy. Next, I study alternative mechanisms for allocating capital in the gig economy: a traditional taxi company-like firm, and a frictionless, on-demand rental market, and how these mechanisms impact ride quantities and prices. Finally, I study the welfare implications of driver quantity restrictions like the one recently passed in New York City.

4.2.1 *Financing costs in the gig economy*

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%. I first ask to what extent this historically low interest rate contributed to the gig economy's growth and employment demographics. In dollar-equivalent terms, the model estimates that the additional cost of finance is roughly \$1,000 per year. I allow \bar{f}^0 , which determines the average cost of finance, to vary between 0 and 6 times its baseline value. $\bar{f}^0 = 0$ captures the case where financing is free. When \bar{f}^0 is three times its current level, this corresponds to the historical cost of finance. Very high levels of \bar{f}^0 correspond to case where financing is effectively unavailable. Each \bar{f}^0 generates an equilibrium price and quantity of ride share rides which flow through to equilibrium driver net income, outside wage, and vehicle distributions.

The model shows 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. What appears on the surface to be a good outcome for drivers—higher wages—only occurs because the lowest-income drivers are excluded from the market.

Figure A.13 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. At historical levels of finance costs, roughly three times current levels, the model estimates there would be roughly 8.5% 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 15% and prices rise to nearly \$30 per hour. The fact that prices increase by nearly 30% when finance is unavailable shows how critical it is to the growth of ride share.

In terms of employment outcomes, changes in the price of finance lead to significant changes in driver income and demographics. Figure A.13 Panel (c) shows yearly net income from full-time ride share driving, defined as the hourly driver income net of Uber commissions, 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 baseline scenario, the net income of a driver is roughly \$15,000 per year, which is just under the national minimum wage of \$7.50 per hour. When excluding the costs of finance, the hourly income is approximately \$8.75, indicating that financing adds a significant cost to driving for Uber.

As financing costs increase, the net income from Uber driving initially decreases because higher financing costs directly increase the total costs of obtaining a car. Surprisingly, as finance costs increase further, the net income of driving for Uber rises sharply, increasing by roughly 15% at historical levels of financing costs and by nearly 100% when finance is unobtainable. This counterintuitive result arises directly from the of the gig economy's capital allocation structure. As financing becomes more expensive, many drivers are unable to obtain cars and therefore exit the market. As these drivers exit, the ride supply shifts inwards and ride prices increase. Higher ride prices mean that the remaining drivers earn higher revenues. Meanwhile, the effects of higher financing costs are muted by selection into who obtains cars and drivers for Uber: the drivers who require financing exit, while the drivers who do not require financing remain. Panel (d) emphasizes how the outside wage of drivers participating in ride share increases as these selection effects take hold. These forces combine to increase driver net income.

These results illustrate two important points about gig economy capital allocation. First, the market structure enables a high degree of flexibility and selection into gig economy employment. In this case, as higher financing costs make gig economy employment less attractive for drivers requiring financing, participation shifts significantly towards individuals who do not require financing. Once financing is very expensive, further cost increases have little effect because the drivers requiring financing have already exited the market. Second, this analysis illustrates the difficulty of using

average gig economy earnings to evaluate the benefits of gig economy employment for low-income individuals. This counterfactual shows a situation in which a policy increases gig economy wages, but only because the policy drives the lowest-income individuals from gig economy employment altogether. Therefore, one can not conclude without further analysis that policies that increase gig economy wages would necessarily benefit current gig economy workers.

4.2.2 *Capital allocation alternatives to the gig economy*

I next study two alternate mechanisms of capital allocation. First, a *traditional* allocation mechanism—a collection of competitive taxi companies. Here, firms obtain cars using finance, hire labor at minimum wage—close to the prevailing after-expenses income of current ride share drivers—and sell rides in a competitive market. I assume the firms’ financing costs are equal to the estimated mean household financing cost, that firms have free entry, and that firms have constant returns to scale technology. These assumptions pin down the equilibrium ride price as the that which makes firm profits equal to zero. Second, I study a *car share* capital allocation mechanism. Here, ride share drivers rent idle cars from other households. As above, I assume that drivers earn minimum wage, meaning that the car owners earn the hourly after-commission ride price, less vehicle expenses, finance expenses, and wage expenses. Note that across these scenarios, I hold fixed the technology of the app, the regulatory environment, and competitive market structure. The only difference is how the cars used for ride share are allocated.

The model shows that the gig economy and car share mechanisms offer significant advantages over the traditional firm structure. 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.

Figure A.14 quantifies this intuition. Panel (a) shows ride share quantities. As compared to a traditional firm, there are roughly 25% more ride share rides in the gig economy. The car share scenario's introduction of a rental market would increase quantities by a further 20%. Panel (b) shows that going from the traditional structure to the gig economy structure decreases the hourly ride price from \$33 per hour to \$22 per hour. A frictionless rental market would decrease prices even further to roughly \$15 per hour. Panel (c) shows driver net income across these three scenarios. Observe that the increases in ride quantities and decreases in prices observed in the gig economy and rental market scenarios are not driven by significant reductions in driver wages. Panel (d) shows the outside income of ride share capital owners. This panel omits the traditional firm scenario, because there, firms, rather than individuals, own cars. In the gig economy scenario, ride share vehicle owners must drive and so they have low outside incomes. In the car share scenario, the average outside income is much higher, because higher income households, who face lower financing costs and value the convenience of owning a car more, own cars and rent them to low-income drivers.

Figure A.15 shows how finance costs impact ride share across the three models of capital allocation. In each panel, the solid line shows the traditional firm structure, the dotted line shows the gig economy structure, and the dashed line shows the car share structure. Panels (a) and (b) show quantities and prices, respectively. As finance costs increase, ride quantities in the traditional firm decline significantly, gig economy prices decrease more slowly and level off, and car share prices decrease very slowly and level off even sooner. The reverse is true of prices, with traditional firm prices increasing linearly, gig economy prices rising somewhat, and car share prices increasing only by a small amount.

These three mechanisms exhibit dramatically different sensitivities to financing costs. This illustrates the mechanisms and consequences of gig economy technology. The traditional firm must own and finance the capital by itself. Consequently, it (and its customers) cannot avoid these costs, leading to price and quantity decreases that occur in lockstep with financing cost increases. In the gig economy, drivers must still own the cars, but there is some flexibility in which individuals own them. When financing becomes more expensive, these cost increases are muted by a substitution towards individuals who do not require finance. This substitution is necessarily incomplete, however, because a substitution towards individuals with better outside employment opportunities, who are costlier ride share drivers. The car share mechanism offers the greatest flexibility because ownership can substitute towards households that do not require financing who rent the cars to low-income drivers regardless of their financing needs.

Panel (c) shows how driver net income varies across these scenarios. The driver net income for gig economy varies as shown earlier in Figure A.13 Panel (c), with incomes for the traditional firm scenario and the car share scenario remaining constant, by assumption. Finally, Panel (d) shows the incomes of ride share car owners across these scenarios. This figure omits the case of the traditional firm, because there firms, rather than households own cars. This figure demonstrates first the large level difference in incomes across the gig economy and car share scenarios, with owners in the gig economy scenario being low-income drivers, and owners in the ride share economy being high-income households renting to drivers. As finance costs increase, outside incomes both rise as low-income, low-liquidity borrowers exit the market. However, the rise is more dramatic in the case of the gig economy scenario because these lower-income drivers are more likely to require financing.

4.2.3 Driver quantity restrictions

Finally, I study limits on the number of ride share drivers. New York City recently implemented a one-year freeze on the number of new ride share licenses. This is an onerous restriction given the fast growth rate of registered Uber drivers in New York: the 2016 year-over-year growth rate was approximately 60%. Limiting quantities pushes the market out of equilibrium. I calculate the ride prices p_d and p_s such that quantity demanded at p_d and supplied at p_s equal the restricted quantity. p_d represents the effective cost to riders, including shadow costs such as increased waiting times. p_s represents the effective income for drivers, including shadow costs such as queuing to obtain an operating license. I do not explicitly model these non-monetary costs.

Figure A.16 shows the results. Panel (a) shows the shadow costs, defined as the difference between p_d and p_s , which increases significantly as the driver restriction becomes larger. Panel (b) shows full-time equivalent driver wages. The figure shows both observed income (solid line) and actual income (dashed line). Observed income represents what an official observing driver net income would see, assuming the transaction price for rides is p_d . Actual income represents what drivers actually receive. A goal of the policy is to increase driver wages. The policy does increase *observed* driver wages significantly, increasing them by nearly 50% for a 25% reduction in the number of drivers from the baseline. The actual wage, after accounting for queuing or other shadow costs, decreases by roughly 6%. Panel (c) shows how driver outside option income changes with the driver restriction. As the driver restriction becomes tighter, the average driver's outcome income falls, because their actual ride share income decreases. This dissuades higher outside income drivers from entering. The decrease of \$200 is relatively modest compared to the impact of higher financing costs or alternate capital allocation mechanisms.

Panel (d) shows welfare losses of drivers (solid) and riders (dashed) in a New York equivalent sized city, in millions of dollars. The most restrictive policy reduces driver welfare by roughly \$15 million per year, and rider welfare by nearly \$300 million per year. The fact that consumers

suffer larger welfare losses than drivers reflects different demand and supply elasticities. Demand is relatively inelastic, while supply is relatively elastic. Supply is elastic because drivers have comparable wage earning alternatives to which they can substitute. The conclusion from this analysis is that while consumers suffer significantly, drivers do not, in large part because of the flexible structure of gig economy employment.

CHAPTER 5

DISCUSSION AND CONCLUSION

I showed that finance has a significant impact on the size and penetration of the gig economy. Access to finance is critical for low-income households to obtain the cars they need to become ride share drivers. 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. I built a structural model to quantify the impact of finance costs, introduced a counterfactual rental market that alleviates financial frictions, and studied the impact of driver quantity restrictions. My model shows that cheap financing increased ride share growth by 18% and decreased prices by 25%. The introduction of a frictionless rental market would increase growth further by 22%, reduce prices by 35%, and lessen the gig economy's reliance on finance. Welfare losses from driver quantity restrictions are large and fall mostly on riders.

5.1 Related literature

5.1.1 Consumer finance

: My paper contributes to the consumer finance literature by highlighting how consumer finance can play a novel role in facilitating production and labor market participation. Studies of consumer finance typically center on how the financial system and financial constraints impact household consumption, durable goods purchases, and savings. For example, Hurst and Stafford (2004), Mian and Sufi (2011), Mian and Sufi (2012b), Ganong and Noel (2018), and many others study aspects of household mortgage borrowing and how this borrowing feeds back into household and aggregate consumption. Other literature has studied alternate financial products and consumer durable goods, such as cars, as in Benmelech et al. (2017) and Mian and Sufi (2012a), or consumer

and payday loans, as in Zinman (2010) or Melzer (2011). Still others study household savings, as in Gennaioli et al. (2015) or financial mistakes and ways that contracts can be designed to alleviate them, e.g., Campbell (2012), Agarwal et al. (2015), Piskorski and Tchistyi (2010), or Jørring (2017). A common thread in much of this literature is that households' underlying motivation for borrowing involves consumption, and by financing expensive acquisitions like houses or cars, they can smooth consumption across time and states of the world.

I extend this literature by studying household finance and financial constraints in a context where finance serves a substantially different purpose. As in other contexts, household finance in the gig economy plays a critical role in allowing low-income households to acquire costly goods. However, rather than obtaining the capital for consumption purposes, these households obtain the capital for production and labor market participation. I show that a lack of finance greatly reduces households' ability to access labor opportunities and skews gig economy employment towards higher-income households who do not require finance. This finding speaks to recent literature evaluating the size and benefits of the financial system, such as Greenwood and Scharfstein (2013) and Philippon (2015). By showing how finance helps low-income households acquire capital to participate in the labor market and increase the growth of the gig economy, my paper shows an important and overlooked channel through which the financial system can improve social welfare.

5.1.2 Corporate finance and productivity

: My paper relates to a long line of research that has emphasized the importance of the financial system in growth and productivity. See, for example, King and Levine (1993), Rajan and Zingales (1996), Kaplan and Zingales (1997), and Porta et al. (1998) highlighting the importance of law in facilitating productivity-improving financial transactions. Khwaja and Mian (2005) pointing to how political interference in financial transactions can reduce productivity. Concerning specific mechanisms, a major driver in finance's importance is in alleviating factor misallocation of the type

studied in Hsieh and Klenow (2009). To the extent that the financial system aids in the reallocation of factors among firms, as shown in Buera et al. (2011), Midrigan and Xu (2014), and Lenzu and Manaresi (2017), remedying these misallocations can lead to large gains in productivity.

My paper extends this literature by documenting how finance is particularly critical for the growth and evolution of disruptive technologies like Uber and Lyft. In doing so, I suggest that finance shapes the boundaries of the gig economy by differentially impacting the penetration of these services across the income distribution. Additionally, I am the first to study how finance transforms household durable consumption goods, such as cars or homes, into productive capital, curing a factor misallocation problem in whether the capital is used for consumption or production. Moreover, I show that finance plays an important role in the reallocation of goods towards those most likely to engage in production in the gig economy and in altering the usage of such goods for those consuming these services. Together, these aspects increase asset utilization in the economy.

5.1.3 Uber, Lyft, and ride share

: Many empirical papers have studied the impact of ride share on consumer 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), and Cook et al. (2018), 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, and of finance in particular on ride share growth. Other papers have explored the question of capital allocation and utilization in the gig economy, typically in a theoretical sense. 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. Often, these papers predict that ride share entry will lead to a smaller capital stock through greater utilization rates. Gong et al. (2017) examine vehicle registrations in China and find that this is not the case. I show that even in the United States, where vehicle ownership was already widespread,

ride share entry increases the size of the capital stock. I find that some theoretical predictions, such as increased utilization, are consistent with the data.

Many papers and policy discussions concerning Uber and Lyft seek to identify where the welfare gains from Uber or Lyft originate. Aside from the frequently flagged better ride provision technology or regulatory arbitrage and increased competition, my work highlights a third reason for ride share's disruptive growth. I show that the structure of the gig economy, which enables the usage of vehicles for both private durable consumption and production, significantly expands the supply of capital used in ride share driving.

5.1.4 Methodology

: Finally, my paper incorporates and extends various consumer demand methodologies from the industrial organization literature, originating with random coefficients demand models from Berry et al. (1995) and Nevo (2001). Other recent papers in consumer finance, e.g., Egan et al. (2017), Benetton (2017), Buchak et al. (2018b), and Buchak et al. (2018a) have adopted these methods. My paper further extends these models by incorporating an explicit financing and labor market choice, allowing me to study the joint relationship between financing, employment, and productivity.

5.2 Discussion

Studying the role of finance in the gig economy illuminates how the gig economy is different than a traditional firm that owns capital and hires labor. In the traditional structure, firms own cars only for production purposes. In the gig economy structure, households own cars both for production, but also for the convenience of having a vehicle. This dual purpose means that there is a larger supply of usable capital stock in the gig economy because households are willing to own cars for non-productive purposes and use them in the gig economy when idle.

This structure, however, still relies on finance because low-income drivers must own the cars that they drive. The fact that ownership is dispersed means that the gig economy is less sensitive to changes in financing costs. When finance costs rise or becomes unavailable, capital ownership can substitute towards households that are less reliant on financing. The ability to substitute is even greater when a frictionless rental market is available, because capital ownership can shift to high-income households that do not rely on financing. These households then rent their cars to low-income drivers who would otherwise depend on financing. Finally, the flexible nature of ride share employment means that welfare losses from driver quantity restrictions fall mostly on riders rather than drivers.

My paper quantifies a novel explanation for the disruptive rise of the gig economy relating to how it redraws firm boundaries. Other explanations for ride share's rise, such as a technological advantage coming from the app's ability to connect drivers with riders and price rides dynamically, or from avoiding taxi cab regulations and increasing competition, ignore a significant technological innovation of the gig economy. This innovation is to make idle consumer durables available for productive use.

My model suggests that this fact alone contributes to a nearly 25% greater gig economy penetration. The gig economy redraws firm boundaries around households. This makes firms very small, but also allows the household-firms' capital to be used for both production and consumption. This is not a new idea, but the gig economy enables this dual usage on a very large scale and in a capital-intensive industry. My paper shows that the gig economy's large-scale blurring of firm boundaries leads to significant gains in productivity. Future research may identify other costs and benefits from how gig economy technologies redraw firm boundaries.


Ride share is still a new technology that impacts a relatively slow-to-adjust capital stock. As the capital stock has more time to adjust, future research may find a shrinking auto stock, declining auto sales, and even greater utilization. Additionally, other large technological changes, such as the adoption of self-driving cars, have the potential to lead to significant changes as low-income

drivers are no longer necessary for ride share driving. The advent of self-driving vehicles will decouple the important link between household income, the value of ride share driving, and the central role of household finance.


APPENDIX A

FIGURES

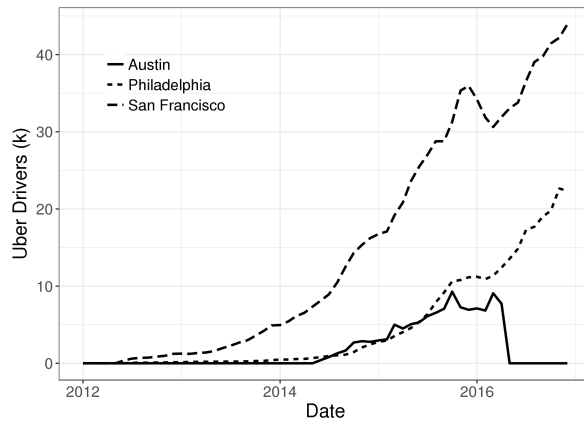
Figure A.1: Timing of ride share entry




(a) New markets entered by month



(b) Total Uber drivers in the United States




(c) Total Uber drivers in representative cities



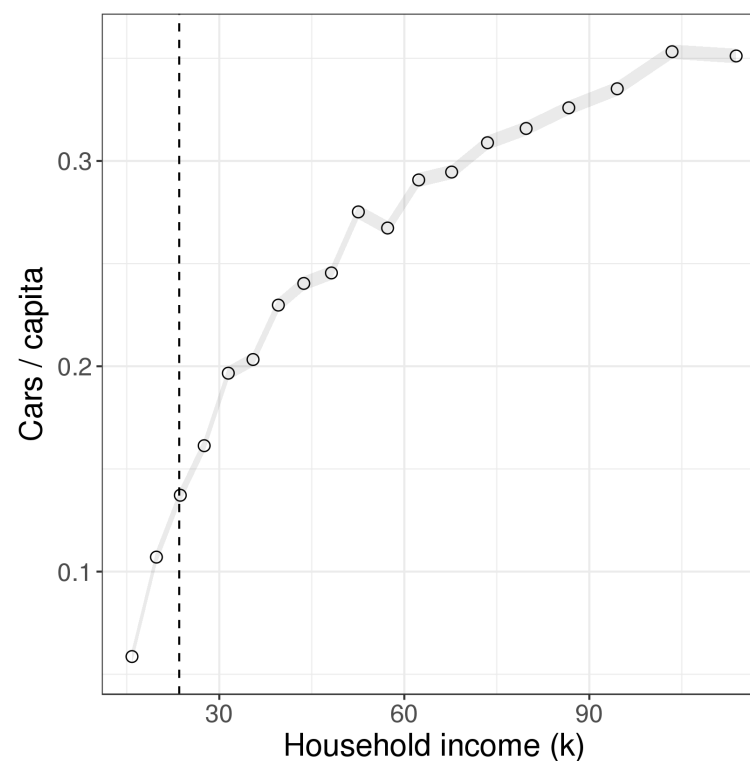
(d) Driver share of population around entry

Figure A.1 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 total number of Uber drivers in representative cities, using Uber data. Panel (d) shows the number of Uber drivers per resident in Uber markets around the time that Uber enters, using Uber data. Gray bars represent 95% confidence intervals.

Figure A.2: Vehicle ownership by household income prior to ride share entry




(a) Vehicles per household versus household income, raw.




(b) Vehicles per household versus household income, CBSA FE.

Figure A.2 shows 2010 vehicle ownership by household income quantile, as calculated in Regression (3.1). Panel (a) shows the raw number of vehicles per household member over eighteen years of age versus household income. Panel (b) shows the number of vehicles per household member over eighteen years of age versus household income with CBSA fixed effects. The vertical dashed line represents the average full-time equivalent ride share driving income of \$23,500. See Mishel (2018). Gray bars represent 95% confidence intervals. Data are from the 2010 ACS.

Figure A.3: Vehicle sales and auto loans after ride share entry



(a) New auto sales



(b) New auto loan originations

Figure A.3 shows the event study for auto sales, Panel (a), and auto loans, Panel (b), as described in Regression (3.2), around ride share entry. 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. The x-axis is time relative to ride share entry in years, with zero representing the quarter in which ride share enters. Gray bars represent 95% confidence intervals with standard errors calculated by bootstrapping across CBSAs.

Figure A.4: Vehicle sales and auto loans after ride share entry by zip code income

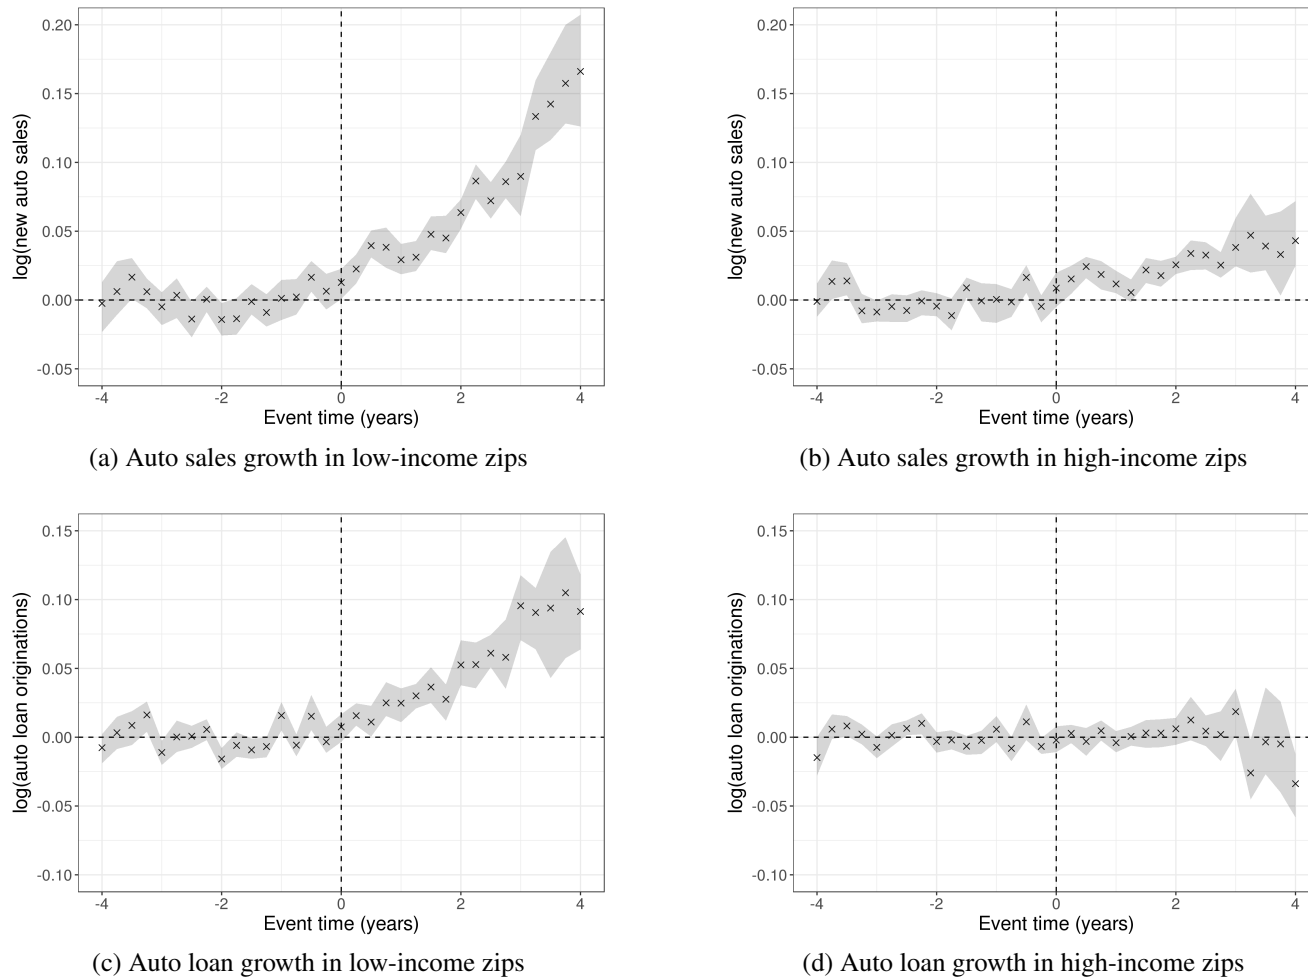




Figure A.4 shows the event study for auto sales, Panels (a) and (b), and auto loans, Panels (c) and (d), as described in Regression (3.2). Data are nationwide at the zip-quarter level from RL Polk (sales) and Equifax (loans). The left-hand side variable is log auto sales or log auto loans. These figures split the sample by relative income of the zip code in the CBSA as of 2010, before ride share entry. Panels (a) and (c) show the effect for zip codes below the CBSA median income and Panels (b) and (d) show the effect for zip codes above the CBSA median income. Gray bars represent 95% confidence intervals, calculated by bootstrapping across CBSAs.


Figure A.5: Tax filings following ride share entry



(a) Tax filings



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



(c) Tax filings with AGI > 25,000

Figure A.5 shows the event study for tax filings as examined in Regression (3.5), 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.

Figure A.6: Vehicle utilization after ride share entry






Figure A.6 shows the event study for vehicle utilization as examined in Regression 3.7. The plot shows the difference between ride share eligible and ride share ineligible utilization in thousands of miles per year within zip codes, plotted against ride share entry event time in years. A vehicle is eligible if it is no more than 15 years old, has four doors, and is a sedan, SUV, or minivan. Gray bars represent 95% confidence intervals. Data comes from the South Carolina DMV vehicle-level registration data set.


Figure A.7: Bankruptcy disclosure and ride share financing




(a) Auto lending probability



(b) Residualized auto lending probability



(c) Cumulative auto lending for constrained borrowers



(d) Cumulative auto lending for unconstrained borrowers

Figure A.7 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) shows the result of Regression 3.8, which 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.

Figure A.8: Experimental design for bankruptcy flag removal

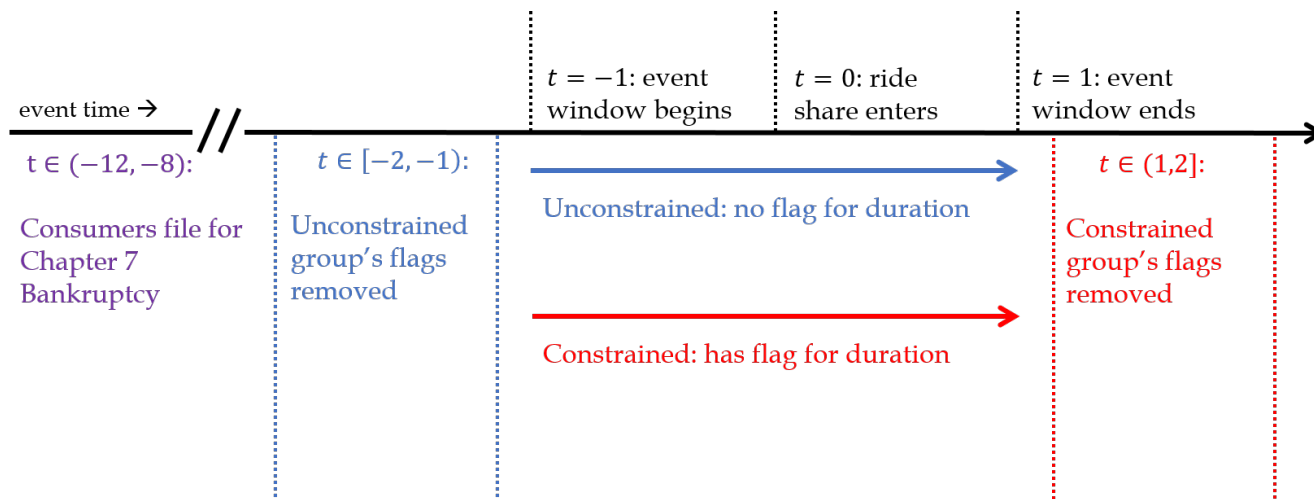


Figure A.8 shows experimental design of the bankruptcy flag removal study. Between 8 and 12 years prior to ride share entry, all consumers in the study file for Chapter 7 bankruptcy. Between 1 and 2 years before ride share entry, the unconstrained group has their bankruptcy flags exogenously removed. 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 A.9: Financial constraints and auto loans after ride share entry

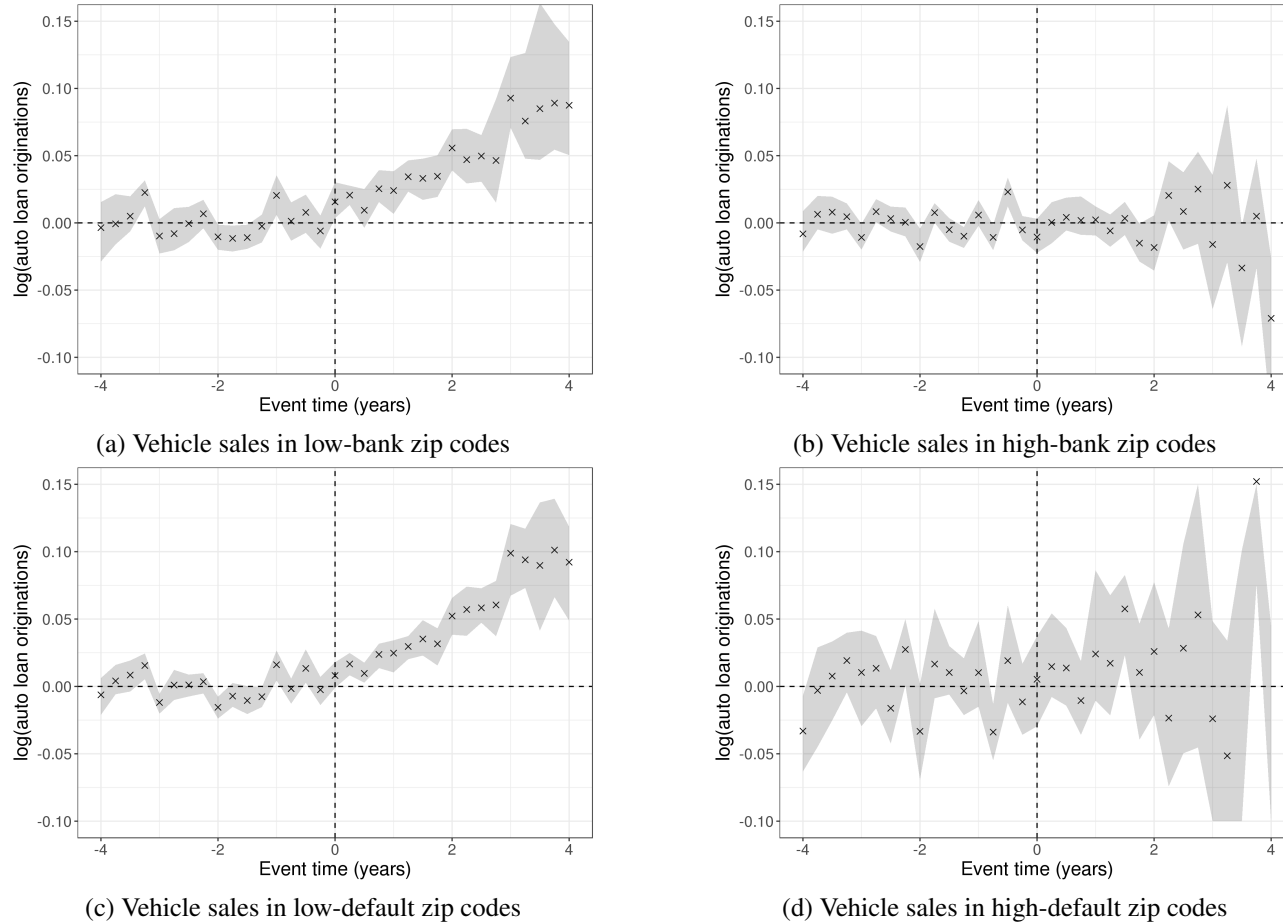


Figure A.9 shows the impact of financial constraints on auto loan origination following ride share entry. The specifications mirror those in Regression (3.2). The left-hand side variable is log auto loan originations, and the sample is split according to financial constraints as measured at the zip code 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.

Figure A.10: Financial constraints and auto sales after ride share entry

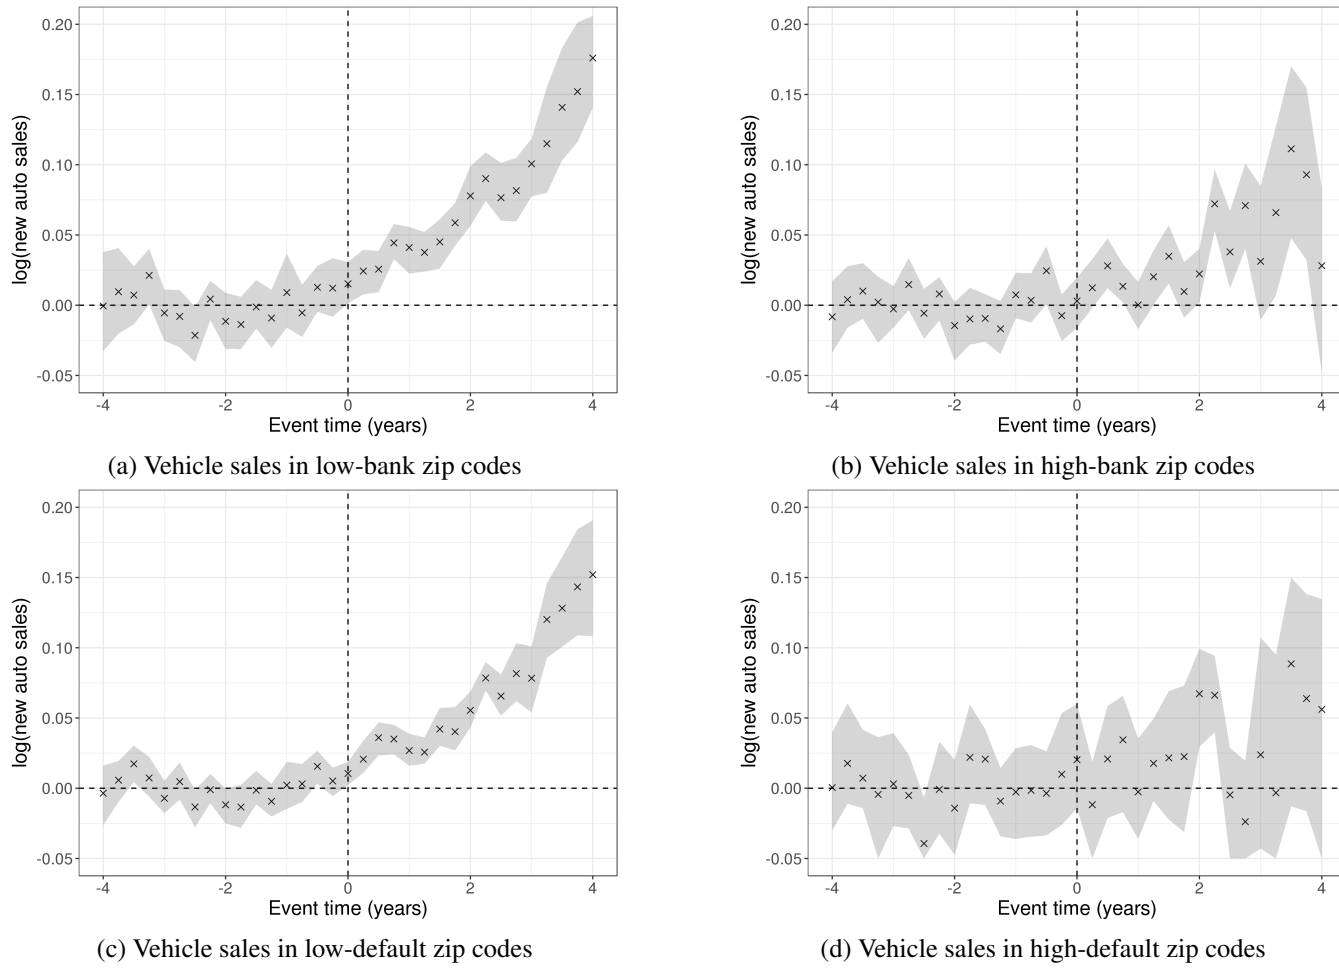


Figure A.10 shows the impact of financial constraints on new auto sales following ride share entry. The specifications mirror those in Regression (3.2). The left-hand side variable is log new auto sales, and the sample is split according to financial constraints as measured at the zip code level. Panels (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.

Figure A.11: Financial constraints and tax filings after ride share entry

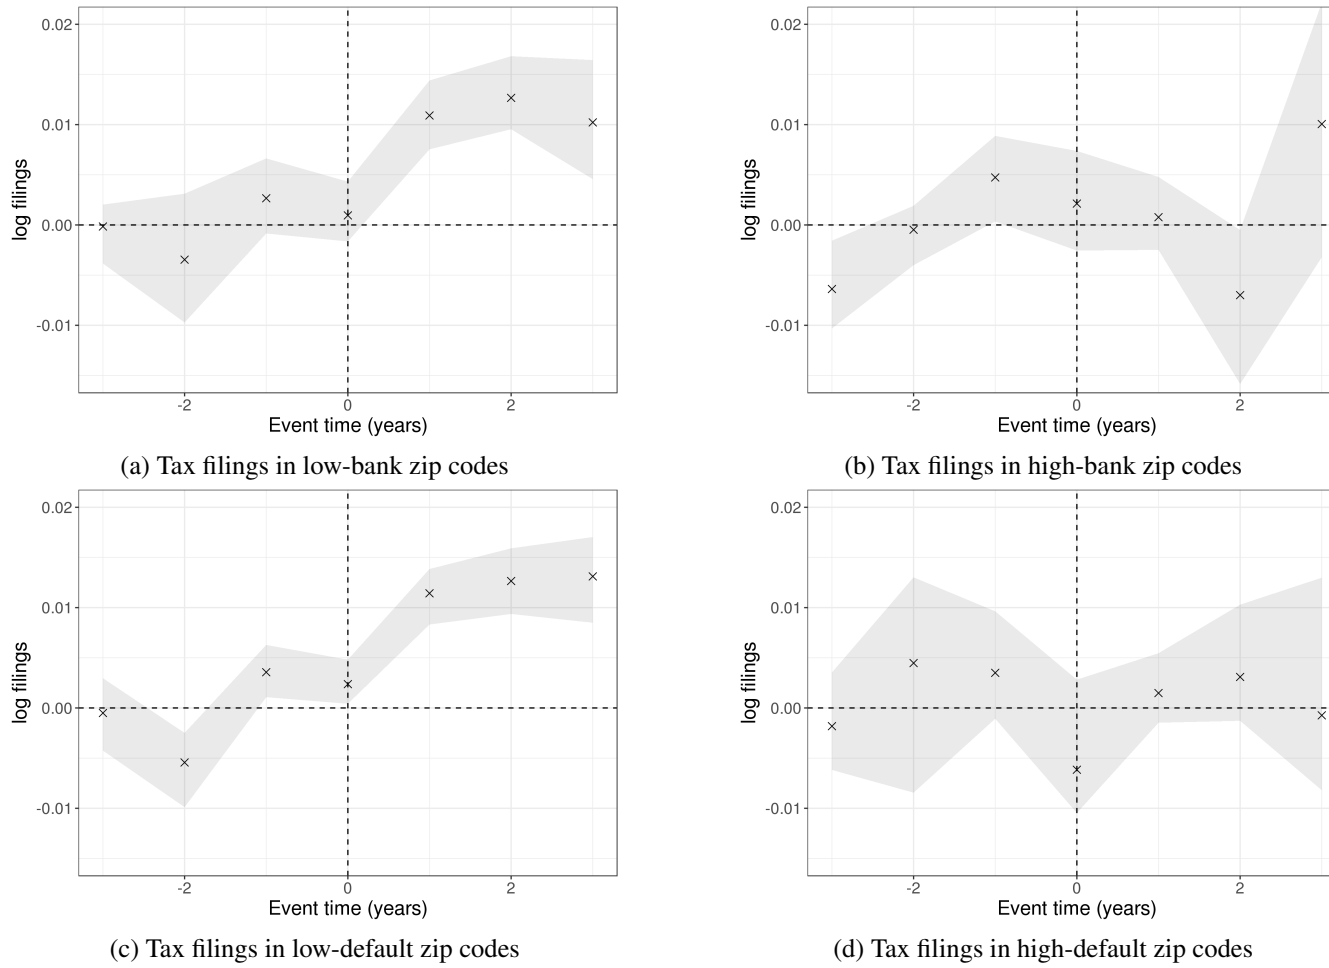


Table A.11 shows the impact of financial constraints on employment following ride share entry. The specifications mirror those in (3.5). The left-hand side variable is log tax filings, and the sample is split according to financial constraints as measured at the zip code level. Panels (a) and (b) split 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.

Figure A.12: Estimation diagnostics: financing and wage

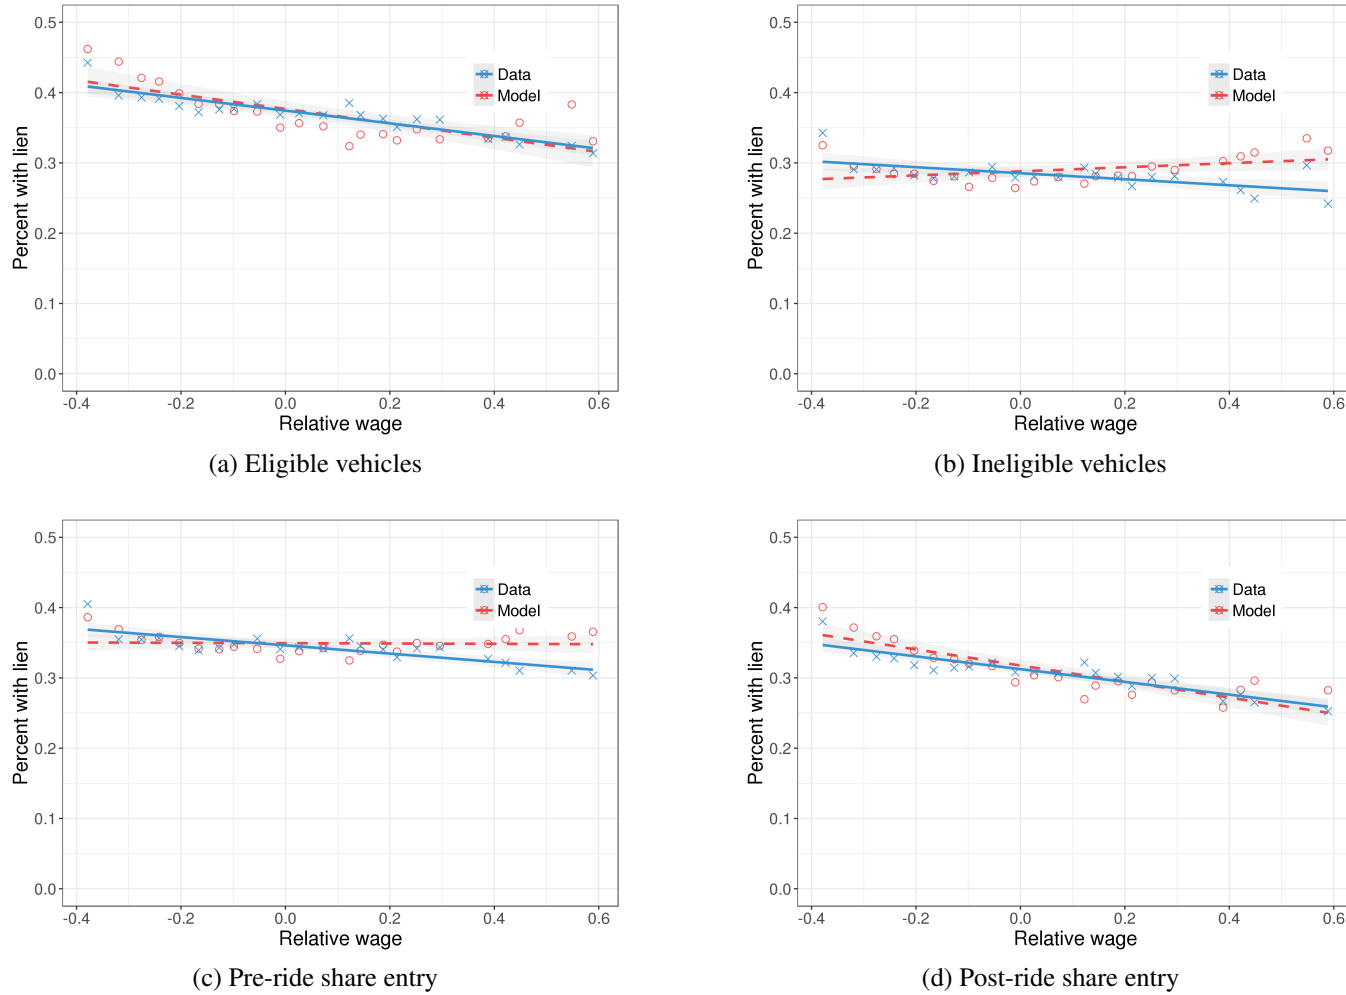


Figure A.12 shows actual (blue solid line and 'x' markers) and model generated (red dashed line and 'o' markers) vehicle financing rates versus individual income. Panel (a) shows financing for ride share-eligible vehicles. Panel (b) shows financing for ride share-ineligible vehicles. Panel (c) shows financing before ride share entry. Panel (d) shows financing after ride share-entry. Gray bars represent 95% confidence intervals.

Figure A.13: Counterfactual costs of finance in the gig economy, prices, quantities, and wages

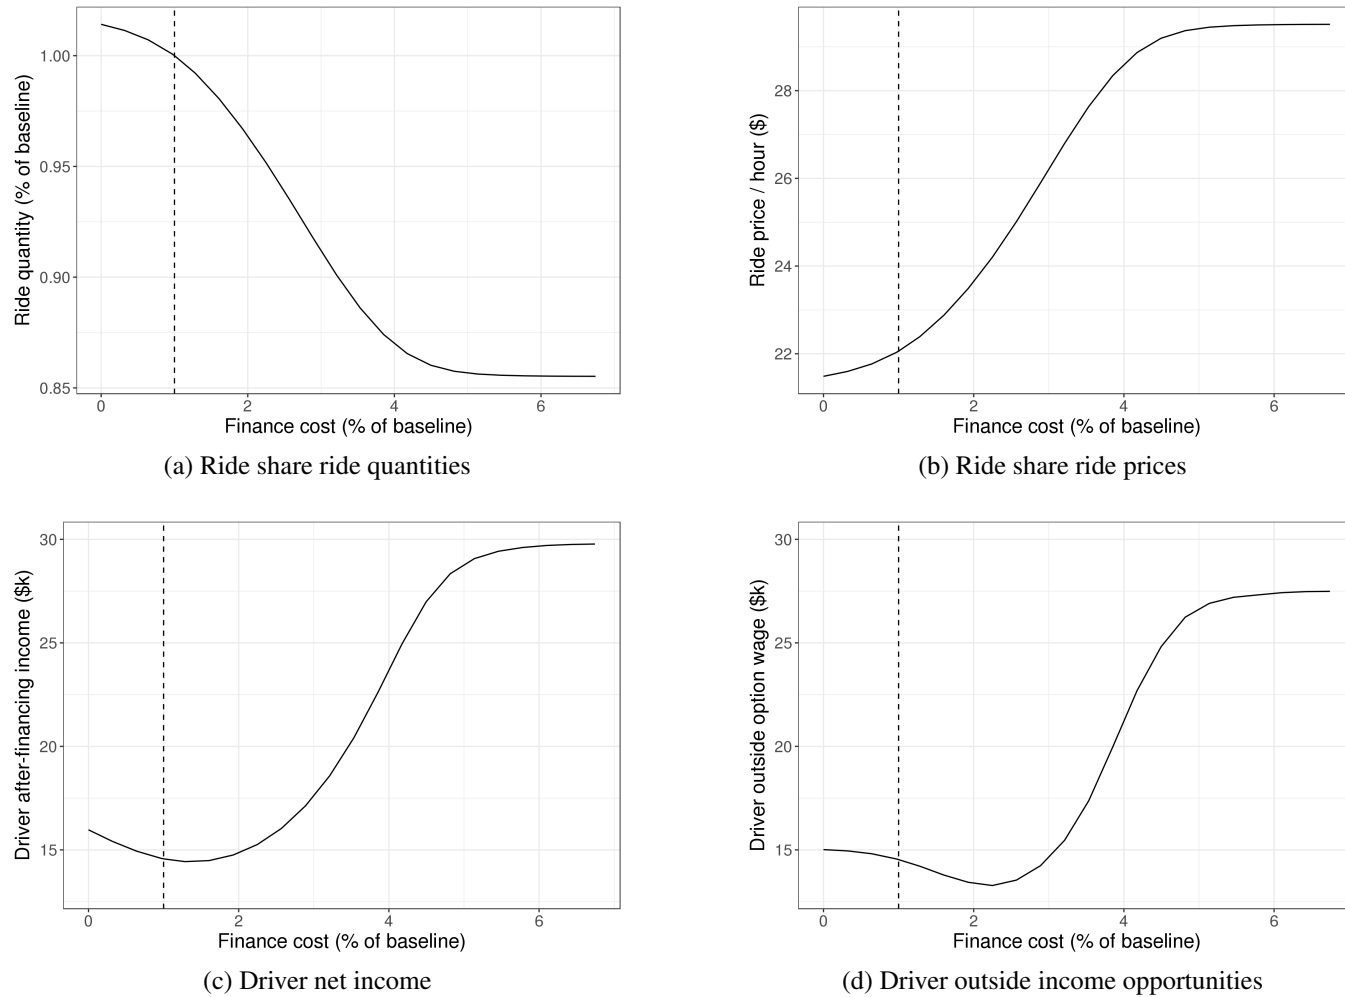


Figure A.13 shows the impact of increasing financing costs. In each panel, the x-axis shows financing costs relative to the baseline estimated value capitalized value of \$14,435, and the vertical dashed line indicates to that value. Panel (a) shows the quantity of ride share rides, normalized to the 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 average yearly outside wage for drivers.

Figure A.14: Counterfactual systems of capital allocation

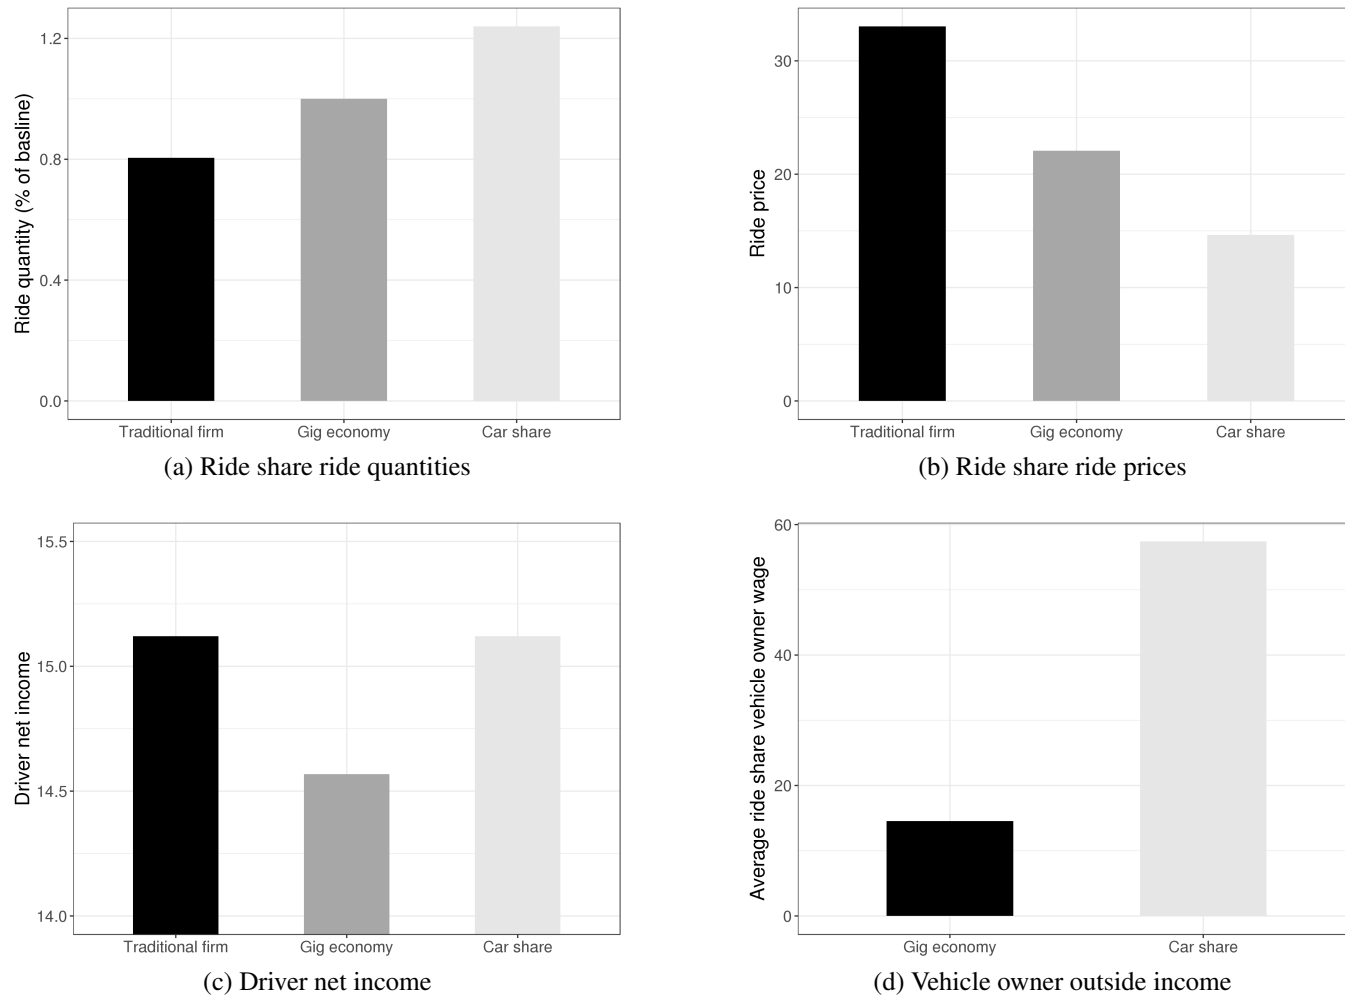


Figure A.14 compares systems of capital allocation. The traditional firm scenario envisions firms owning financed cars and hiring drivers at minimum wage. The gig economy scenario is the current structure, where drivers work as independent contractors and own their cars. The car share scenario envisions 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 outside income of ride share vehicle owners.

Figure A.15: Finance costs and counterfactual systems of capital allocation

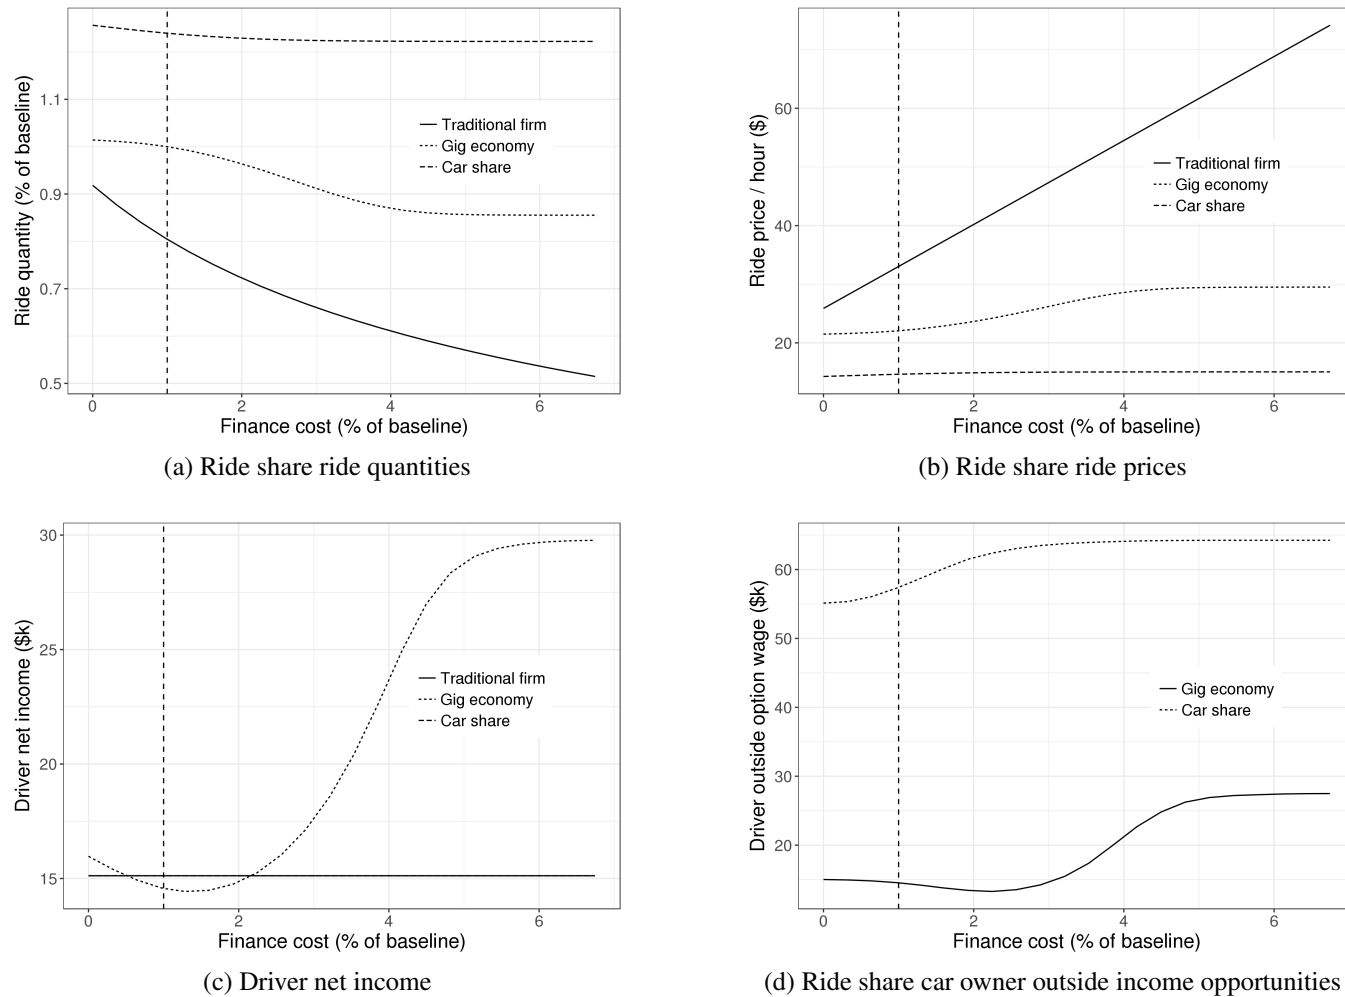
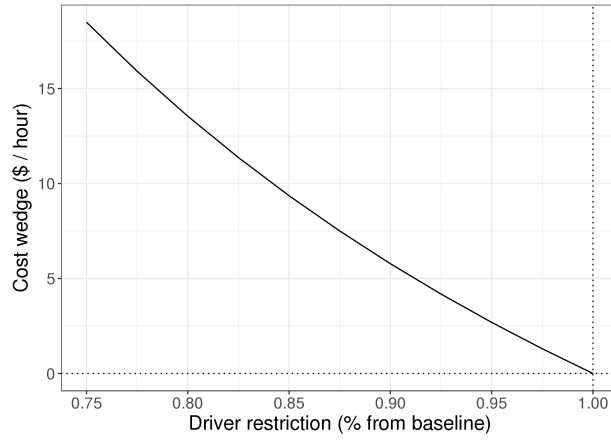
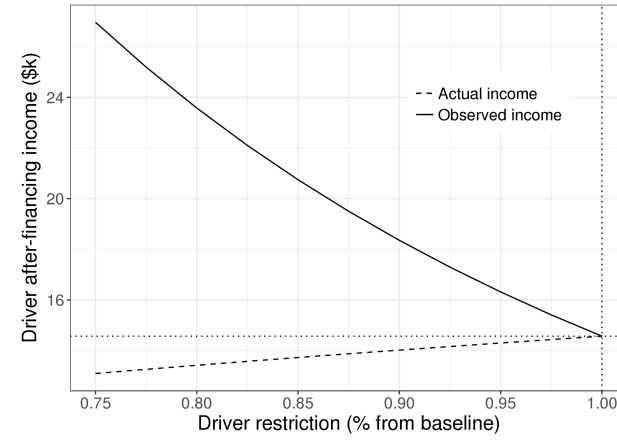


Figure A.15 shows how finance impacts ride share outcomes under different systems of capital allocation. A traditional firm owns financed cars and hires labor. The gig economy scenario is the current structure where drivers work as independent contractors and own cars. The car share scenario envisions 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 outside income of ride share vehicle owners.


Figure A.16: Counterfactual limits on drivers



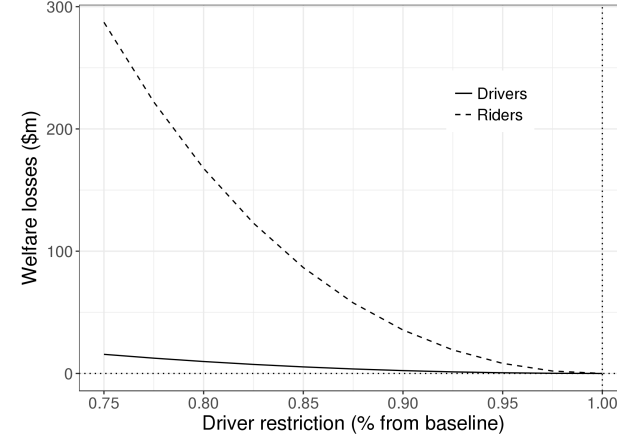
(a) Ride shadow cost



(b) Driver observed and actual incomes



(c) Driver outside wage



(d) Consumer and driver welfare losses

Figure A.16 shows how limiting the number of full-time equivalent Uber drivers impacts ride share outcomes. Driver numbers, on the x-axis, are relative to the baseline, unconstrained case. Panel (a) the shadow cost of the policy, i.e., the wedge between transacted price and the price that drivers receive. Panel (b) shows observed and actual driver net income. Panel (c) shows average driver outside option income. Panel (d) shows the welfare losses for consumers and drivers.

APPENDIX B

TABLES

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

Panel A: New auto sales, nationwide, RL Polk data			
	<i>Dependent variable:</i>		
	Log sales		
	(1)	(2)	(3)
Post	0.016*** (0.002)	0.006*** (0.002)	0.004 (0.003)
Post \times Low income	-	0.020*** (0.004)	-
Post \times High transport share	-	-	0.025*** (0.004)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr \times Low income FE	N	Y	N
Qtr \times High transport FE	N	N	Y
Observations	244,153	244,153	244,153
R ²	0.971	0.972	0.972
<i>Note:</i> *p<0.1; **p<0.05; ***p<0.01			

Panel B: New auto loan originations, nationwide, Equifax data			
	<i>Dependent variable:</i>		
	Log new originations		
	(1)	(2)	(3)
Post	0.010*** (0.002)	-0.001 (0.002)	-0.002 (0.002)
Post \times Low wage	-	0.021*** (0.003)	-
Post \times High transport share	-	-	0.024*** (0.003)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr \times Low Wage FE	N	Y	N
Qtr \times High transport FE	N	N	Y
Observations	244,153	244,153	244,153
R ²	0.979	0.979	0.979
<i>Note:</i> *p<0.1; **p<0.05; ***p<0.01			

Table B.1 shows the results of Regressions (3.3) and (3.4). 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. The regressions are on the zip-quarter level, and only zip codes that eventually see entry are in the regression. In each panel, 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 in parentheses are clustered at the CBSA-quarter level.

Table B.2: Employment after ride share entry

	<i>Dependent variable:</i>		
	log filings		
	(1)	(2)	(3)
Post	0.007*** (0.002)	0.002 (0.003)	- -
Post \times (AGI \leq 25k)	- -	0.010** (0.004)	0.010** (0.004)
(AGI $<$ 25k) \times Year FE	Y	Y	Y
(AGI $<$ 25k) \times Zip FE	Y	Y	Y
Zip \times Year FE	N	N	Y
Observations	155,246	155,246	155,246
R ²	0.996	0.996	0.998
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table B.2 shows the result of regression (3.6) 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. *AGI \leq 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.

Table B.3: Vehicle utilization after ride share entry

	<i>Dependent variable:</i>					
	Miles per year (k)					
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.175 (0.160)	−0.156 (0.282)	−1.033** (0.444)	- -	0.299 (0.659)	- -
Post × Low income	- -	0.381 (0.279)	- -	- -	−1.518** (0.605)	- -
Post × Eligible	- -	- -	1.350*** (0.469)	1.186** (0.543)	−0.561 (0.714)	−0.922 (0.905)
Post × Low income × Eligible	- -	- -	- -	- -	2.185*** (0.663)	2.372*** (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

Note:

*p<0.1; **p<0.05; ***p<0.01

Table B.3 shows the result of Regression (3.7) using DMV registration data from South Carolina between 2010 and 2016. The left-hand side variable is the number of miles driven per year in thousands during the ownership spell. The observations are at the individual vehicle-ownership level. *Post* is an indicator for whether ride share has entered. *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, i.e., no older than 15 years old and a four-door sedan, minivan, or SUV. All columns include zip × *Eligible* and quarter × *Eligible* fixed effects. Columns (4) and (6) include zip × quarter fixed effects. Standard errors, in parentheses, are clustered at the CBSA-quarter level.

Table B.4: Bankruptcy disclosure and ride share financing**Panel A: Bankruptcy flag and auto originations**

	<i>Dependent variable:</i>				
	P(auto loan) (%)				
	(1)	(2)	(3)	(4)	(5)
Window (years)	±0.25	±0.50	±1.00	±1.50	±2.50
≥ 10 years	0.132*** (0.032)	0.149*** (0.021)	0.132*** (0.015)	0.136*** (0.012)	0.115*** (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

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

Panel B: Bankruptcy flag removal and ride share entry

	<i>Dependent variable:</i>					
	P(auto loan) (%)					
	Window = ± 1 year			Window = ±0.50 years		
	(1)	(2)	(3)	(4)	(5)	(6)
Post _{zt}	0.085 (0.059)	0.243*** (0.083)	- -	0.093 (0.081)	0.255** (0.109)	- -
Post _{zt} × Constrained	- -	-0.316** (0.128)	-0.317** (0.142)	- -	-0.310* (0.171)	-0.326* (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

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

Table B.4 shows the impact of bankruptcy disclosure on auto lending and ride share entry. Panel (A) shows the result of Regression (3.9), studying the impact of flag removal on the level of auto lending. Columns (1)–(4) vary the observation window around the ten-year removal window, expanding the window from ±0.25, 1.00, 1.50, and 2.50, respectively. ≥10 years is an indicator for whether the filing occurred more than 10 years from the observation date. Panel (B) shows the result of Regression (3.10). 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.

Table B.5: Financial constraints and ride share entry**Panel A: Financial constraints and auto lending following ride share entry**

	<i>Dependent variable:</i>	
	log new originations	
	(1)	(2)
Post	0.021*** (0.005)	0.024*** (0.005)
Post \times High bank share	-0.080*** (0.017)	- -
Post \times High default share	- -	-0.078*** (0.019)
Zip FE	Y	Y
Credit \times Income \times Time FE	Y	Y
Post \times 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

	<i>Dependent variable:</i>			
	log sales		log tax filings	
	(1)	(2)	(3)	(4)
Post	0.022*** (0.006)	0.025*** (0.006)	0.009*** (0.002)	0.009*** (0.002)
Post \times High bank share	-0.038** (0.017)	- -	-0.008** (0.004)	- -
Post \times High default share	- -	-0.048*** (0.014)	- -	-0.010** (0.004)
Zip FE	Y	Y	Y	Y
Credit \times Income \times Time FE	Y	Y	Y	Y
Post \times 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

Table B.5 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 \times High income* controls, not shown in this table. Additionally, all columns include zip fixed effects, and *High bank* or *High default \times Low income* \times 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.

Table B.6: Model parameters**Panel A: Estimated parameters**

Parameter	Description	$\bar{\beta}$	Π_{wage}	Σ
β_0^{pre}	Convenience value (pre)	0.265	-1.755	0.548
β_0^{post}	Convenience value (post)	1.609	-1.163	0.548
m^{pre}	Utilization (k miles) pre	11.552	-	1.969
m^{post}	Utilization (k miles) pos	12.347	-	1.969
$\log \alpha^{pre}$	(Log price sensitivity (pre)	-1.534	-	0.075
$\log \alpha^{post}$	(Log price sensitivity (post)	-1.534	-	0.075
β_e^{pre}	Eligibility value (pre)	2.513	-	3.109
β_e^{post}	Eligibility value (post)	1.271	-	3.109
β_{hp}^{pre}	HP value (pre)	0.805	-	0.322
β_{hp}^{post}	HP value (post)	2.529	-	0.322
β_{mpg}^{pre}	MPG value (pre)	-0.737	-	0.477
β_{mpg}^{post}	MPG value (post)	0.718	-	0.477
f_0^{pre}	Financing cost (pre)	1.897	-1.514	0.739
f_0^{post}	Financing cost (post)	3.114	-1.514	0.739
l^{pre}	Liquidity (pre)	10.826	1.492	0.300
l^{post}	Liquidity (post)	10.415	1.492	0.300
γ	Drive utility	-3.306	-	1.400

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)

Table B.6 shows the model parameters. Panel (a) 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. Panel (b) shows parameters taken from the literature and their sources.

APPENDIX C

SUPPORTING ANALYSES

This appendix contains supplemental material. Section C.1 details the data sources used in the paper, presents summary statistics, and compares DMV registrations to new auto sales. Section C.2 shows regressions concerning the intensive and extensive margins of ride share entry. Section C.3 studies the ex-ante distribution of vehicle ownership in the economy.

Section C.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. Table C.6 contains the analysis of auto loan performance following ride share entry. Tables C.7 through C.11 show placebo tests of all major reduced form results, with the timing of entry preserved but assignment randomized, to ensure that there is no effect, and that the controls and fixed effects properly absorb systematic trends in the data.

C.1 Data

My paper brings together a number of datasets described here. Summary statistics are shown in Table C.1.

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 C.2 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.

Table C.1: Summary statistics

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 C.1 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.

Table C.2: New registrations and new sales comparison

	<i>Dependent variable:</i>					
	Sales		Log sales		Sales / capita	
	(1)	(2)	(3)	(4)	(5)	(6)
New registrations	0.353*** (0.002)	0.339*** (0.004)	- -	- -	- -	- -
Log new registrations	- -	- -	0.989*** (0.002)	0.261*** (0.009)	- -	- -
New registrations per capita	- -	- -	- -	- -	0.400*** (0.004)	0.416*** (0.004)
Constant	-7.050*** (0.821)	- -	-1.214*** (0.010)	- -	-0.002*** (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

Table C.2 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.

C.2 The timing and location of ride share entry

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

$$\begin{aligned} Entered_c &= X'_c\beta + \epsilon_c \\ YearsToEntry_c &= X'_c\beta + \epsilon_c \end{aligned} \tag{C.1}$$

$Entered_c$ 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;¹ the change variables are calculated as three year changes. Table C.3 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.²

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

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

2. 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.

and Lyft enter cities with large markets: large cities where consumers have the mobile broadband access necessary to use the ride share applications.

Table C.3: Timing of ride share entry

	<i>Dependent variable:</i>			
	Entered		Years to entry	
	(1)	(2)	(3)	(4)
log population	0.229*** (0.021)	0.235*** (0.025)	−0.799*** (0.071)	−0.739*** (0.088)
Δ population	1.450 (1.261)	0.815 (1.436)	−4.209 (4.508)	−0.498 (5.173)
HH with mobile broadband (%)	0.754** (0.332)	0.925** (0.377)	−5.905*** (1.202)	−5.150*** (1.457)
Δ % HH with mobile broadband	0.064 (0.319)	0.159 (0.341)	−2.486* (1.412)	−2.861* (1.500)
HH with vehicles (%)	0.364 (1.000)	1.266 (1.073)	−0.619 (3.285)	−2.092 (3.580)
Δ % HH with vehicles	−1.322 (1.326)	−0.678 (1.447)	8.664 (5.872)	7.715 (6.216)
Bank share of auto financing	0.009 (0.186)	0.098 (0.218)	−0.955 (0.691)	−0.204 (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 C.3 shows the determinants of ride share entry as specified in Regression C.1. 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* is 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.

C.3 Ex-ante determinants of vehicle ownership

Figure A.2 in the body of the paper shows ex-ante car ownership per household versus household income. To study a richer set of determinants of vehicle ownership and sales, I run the following zip-level regression:

$$\begin{aligned} Y_z &= X'_z\beta_1 + X'_c\beta_2 + \epsilon_z \\ Y_z &= X'_z\beta_1 + \gamma_c + \epsilon_z \end{aligned} \tag{C.2}$$

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.³ The specification 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 C.4 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.

3. 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 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.

Table C.4: Vehicle stocks and sales before ride share entry

	<i>Dependent variable:</i>			
	Cars per household		Sales per capita	
	(1)	(2)	(3)	(4)
Unemployment rate (zip)	−0.015*** (0.001)	−0.016*** (0.001)	−0.016*** (0.002)	−0.014*** (0.003)
Log income (zip)	0.240*** (0.012)	0.244*** (0.012)	0.358*** (0.047)	0.236*** (0.052)
Average commute time (zip)	0.014*** (0.0003)	0.014*** (0.0003)	−0.005*** (0.001)	−0.006*** (0.002)
Log income (CBSA)	−0.003 (0.017)	-	−0.279*** (0.069)	-
HH with high-speed Internet (CBSA)	−0.472*** (0.025)	-	0.377*** (0.104)	-
HH with mobile broadband (CBSA)	0.725*** (0.039)	-	−1.106*** (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

Table C.4 shows the result of Regression C.2. 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 C.3. 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.

C.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.⁴

I begin with an event studying testing how new registrations of eligible vehicles in low wage zip

4. 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.

codes respond differentially to ride share entry:

$$\log Regs_{mezt} = \sum_{\tau=-4}^4 \beta_{\tau} I(t - ET_z = \tau) \times Low\ income_z \times Eligible_{me} + \gamma_{tme} + \gamma_{zme} + \epsilon_{mezt} \quad (C.3)$$

$Eligible_{me}$ is a zero-one indicator for whether the vehicle is eligible. $Low\ income_z$ is a zero-one 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 C.1 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:

$$\begin{aligned}
\log Regs_{mezt} = & \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Eligible_{me} + \beta_3 Post_{zt} \times Low\ income_z \\
& + \beta_4 Post_{zt} \times Eligible_{me} \times Low\ income_z + t \times Eligible_{me} \times Low\ income_z \\
& + \gamma_{mez} + \gamma_{met} + \gamma_{zt} + \epsilon_{mezt}
\end{aligned}
\tag{C.4}$$

Variables in the regression are as defined previously in equations (3.4) and (C.3). 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 (C.4) 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 C.5 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 B.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 C.8 Panel A and show no significance.

DMV liens data: The preceding analysis used nationwide auto lending data at the zip quarter level. 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:

$$\begin{aligned} HasLien_{vzt} = & \beta_1 Post_{zt} + \beta_2 Post_{zt} \times Eligible_v + \beta_3 Post_{zt} \times Low\ income_z \\ & + \beta_4 Post_{zt} \times Eligible_v \times Low\ income_z + t \times Eligible_v + FE + \epsilon_{vzt} \end{aligned} \quad (C.5)$$

$HasLien_{vzt}$ 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, intended to capture varying trends in likelihood of obtaining financing that varies by vehicle eligibility. Other terms are as described in equation (3.7). 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 C.5 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 a small decrease in the likelihood of a vehicle being financed of roughly

0.20%. Column (2) adds the post times low-income interaction with the same fixed effects, and finds a significant differential effect for low-income zip codes. Vehicles registered in low-income zip codes are 1% more likely to have liens following ride share entry as compared to high-income zip codes.

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 C.8, Panel B, shows no effects.

Figure C.1: New vehicle registrations following ride share entry

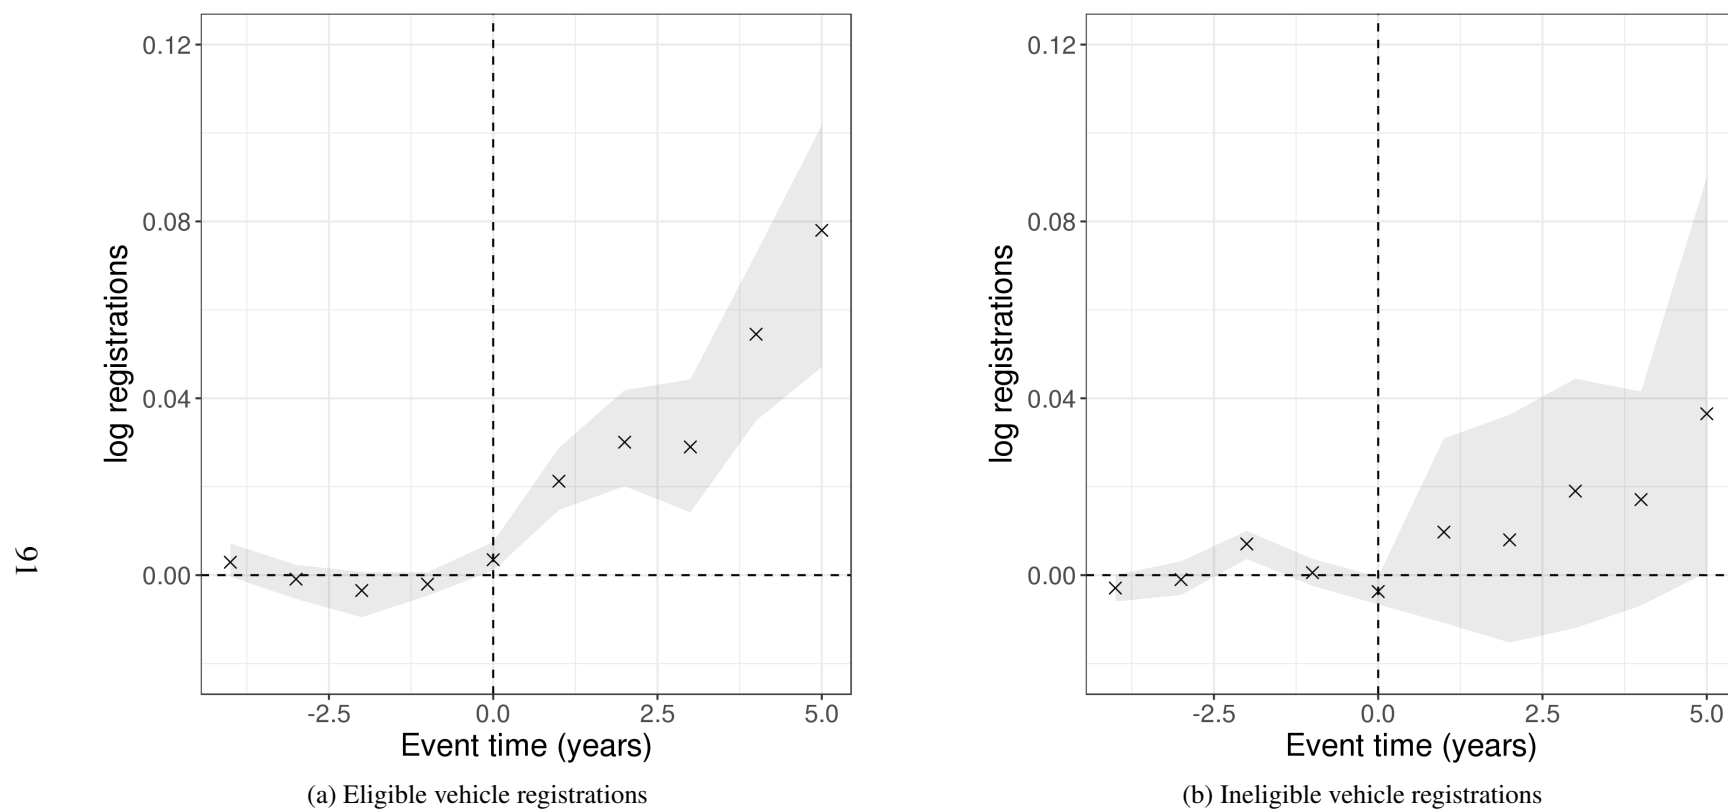


Figure C.1 shows the event study for vehicle registrations as examined in Regression (C.3), which uses DMV data available for Washington, South Carolina, and Indiana. The left-hand side variable is log auto registrations in the zip code. Panel (a) shows the difference-in-difference coefficient for ride-share 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.

Table C.5: DMV registrations and liens**Panel A: New auto registrations, select states, DMV data**

	Log new registrations			
	(1)	(2)	(3)	(4)
Post	−0.0004 (0.016)	−0.005 (0.017)	0.007 (0.014)	-
Post × Eligible	-	-	−0.018*** (0.007)	−0.015** (0.007)
Post × Low income	-	0.010* (0.006)	−0.003 (0.005)	-
Post × Low income × Eligible	-	-	0.019*** (0.005)	0.019*** (0.005)
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

	Has lien			
	(1)	(2)	(3)	(4)
Post	−0.002*** (0.001)	−0.011*** (0.001)	−0.005*** (0.001)	-
Post × Eligible	-	-	−0.007*** (0.002)	−0.001 (0.002)
Post × Low income	-	0.009*** (0.001)	−0.006*** (0.002)	-
Post × Low income × Eligible	-	-	0.019*** (0.002)	0.012*** (0.002)
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	23,897,450	17,724,938	17,724,938	17,724,938
R ²	0.153	0.162	0.073	0.165

Note:

*p<0.1; **p<0.05; ***p<0.01

Table C.5 Panel A shows the results of Regression (C.4) using DMV registration data from select states between 2010 and 2016. The left-hand side variable is log new registrations. *Post* is an indicator taking the value 1 after ride share enters. *Eligible* is an indicator for whether the vehicle is ride share eligible. *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 linear time trends by high- and low-income zip times *Eligible*. Column (4) contains zip-quarter fixed effects. Panel B shows the results of Regression (C.5) at the registration level. The left-hand side is an indicator for whether there is a lien. Columns (1)–(3) contain zip fixed effects, time times low and high income fixed effects, make-model fixed effects, and linear time trends for eligible and ineligible vehicles. Column (4) has zip × time fixed effects. Standard errors in parentheses are clustered at the CBSA-quarter level.

C.5 Loan performance after ride share entry

I study the impact of ride share entry on loan performance. To do so, I repeat specification (3.3), where the outcome of interest is loan performance. In particular, I calculate 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 C.6 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.

Table C.6: Auto loan performance following ride share entry

	<i>Dependent variable:</i>		
	% In Default		
	(1)	(2)	(3)
Post	−0.0003** (0.0001)	−0.001*** (0.0001)	−0.0005*** (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	244,153	244,153	244,153
R ²	0.590	0.593	0.591

Note:

*p<0.1; **p<0.05; ***p<0.01

Table C.6 shows how auto loan performance changes in response to ride share entry. The specification exactly mirrors that in Regression (3.4) 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. Data are from Equifax between 2010 and 2016. Standard errors, clustered at the CBSA-quarter level, are in parentheses.

C.6 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 $\pm 50\%$ around its current value. That is, I consider it taking the values $\{0.29, 0.57, 0.86\}$. Figure C.2 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 C.3 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 C.4 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 C.5 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.

Figure C.2: Counterfactual costs of finance, elasticity sensitivity

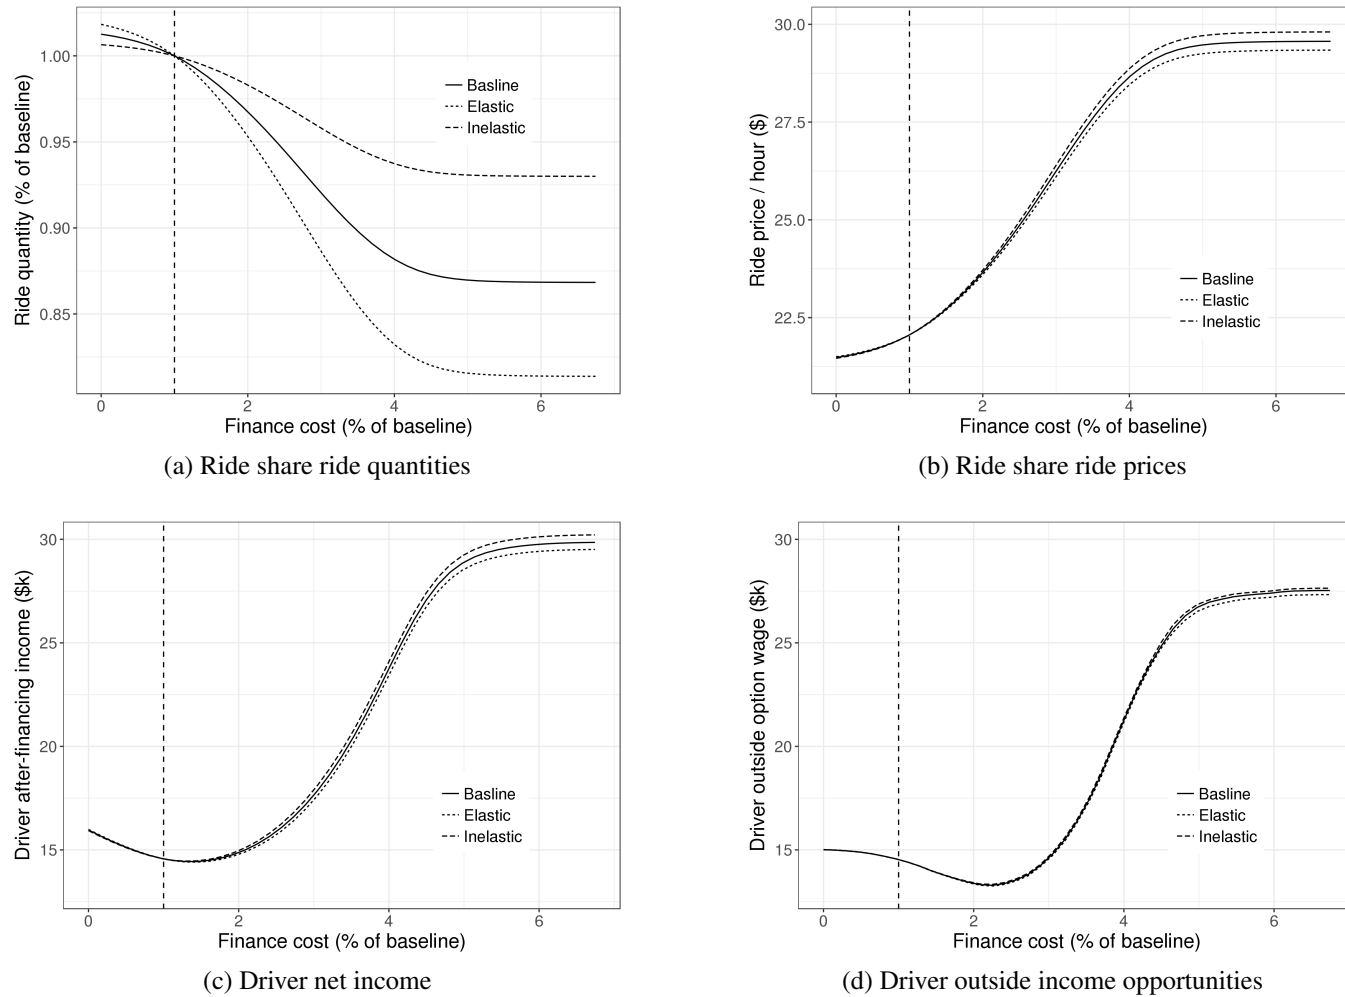


Figure C.2 shows the impact of increasing financing costs for low (0.29), baseline (0.57), and high (0.86) levels of demand elasticity. In each panel, the x-axis shows financing costs relative to the baseline estimated value capitalized value of \$14,435, and the vertical dashed line indicates to that 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 average yearly outside wage for drivers.

Figure C.3: Counterfactual systems of capital allocation, elasticity sensitivity

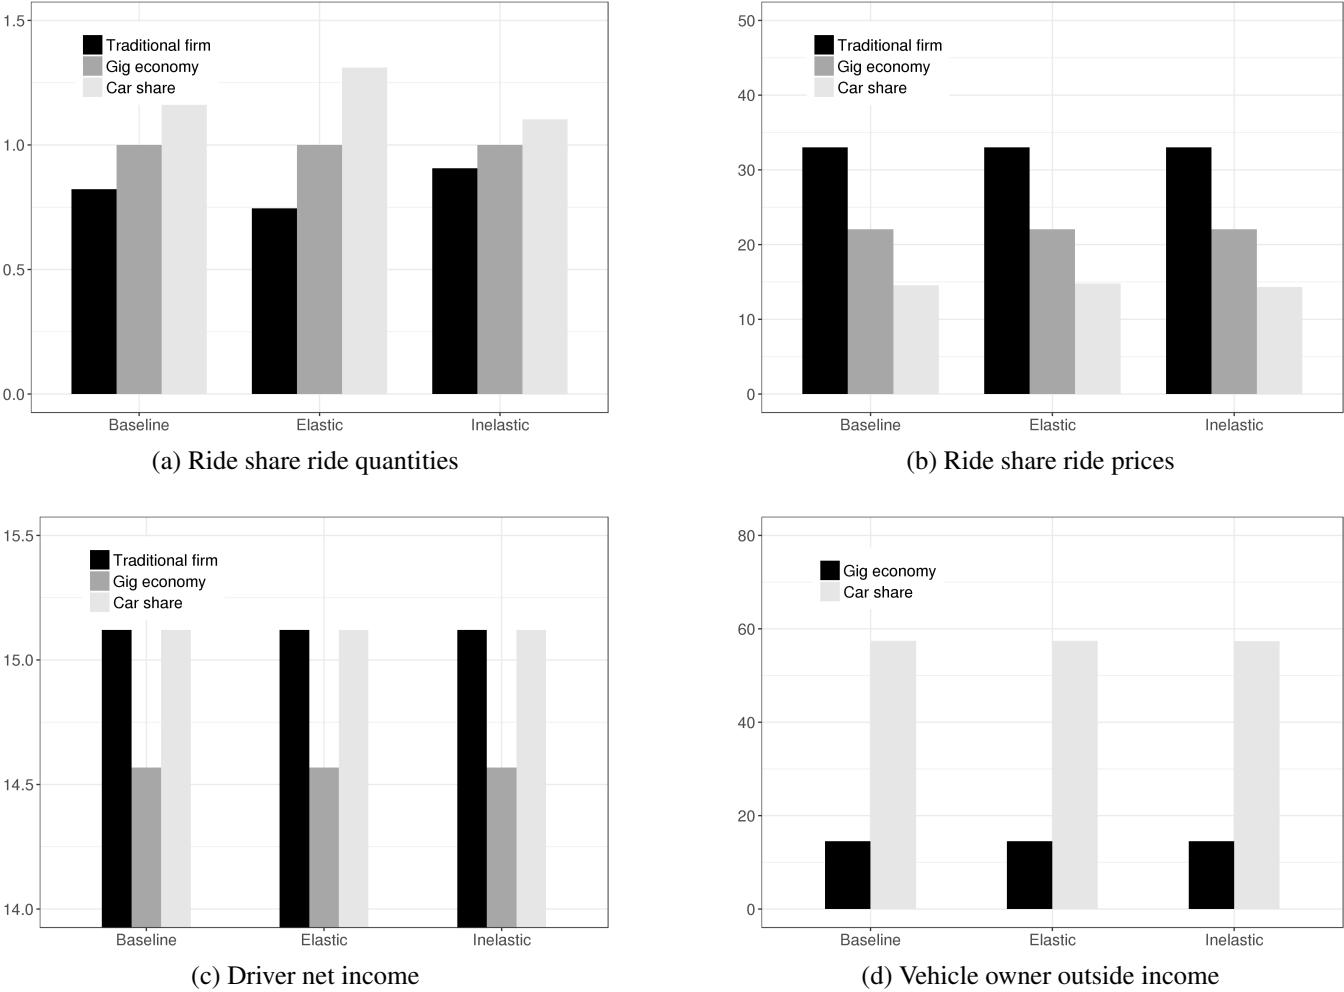


Figure C.3 compares systems of capital allocation for low (0.29), baseline (0.57), and high (0.86) levels of demand elasticity. The traditional firm scenario envisions firms owning financed cars and hiring drivers at minimum wage. The gig economy scenario is the current structure, where drivers work 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 outside income of ride share vehicle owners.

Figure C.4: Counterfactual costs of finance, commission sensitivity

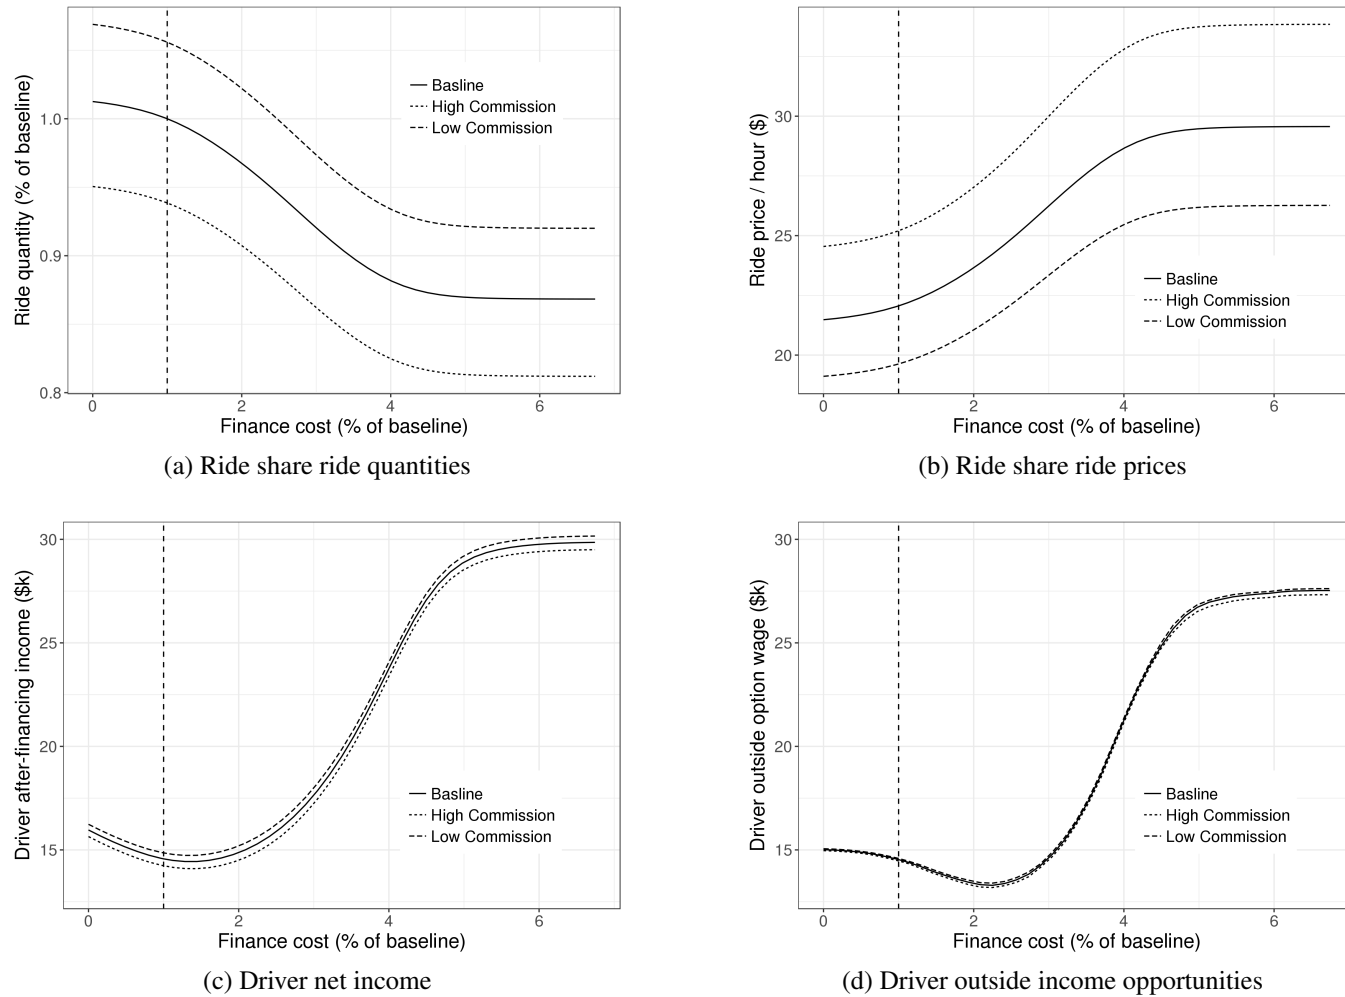


Figure C.4 shows the impact of increasing financing costs for low (15%), baseline (25%), and high (35%) ride share commissions. In each panel, the x-axis shows financing costs relative to the baseline estimated value capitalized value of \$14,435, and the vertical dashed line indicates to that value. 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 average yearly outside wage for drivers.

Figure C.5: Counterfactual systems of capital allocation, elasticity sensitivity

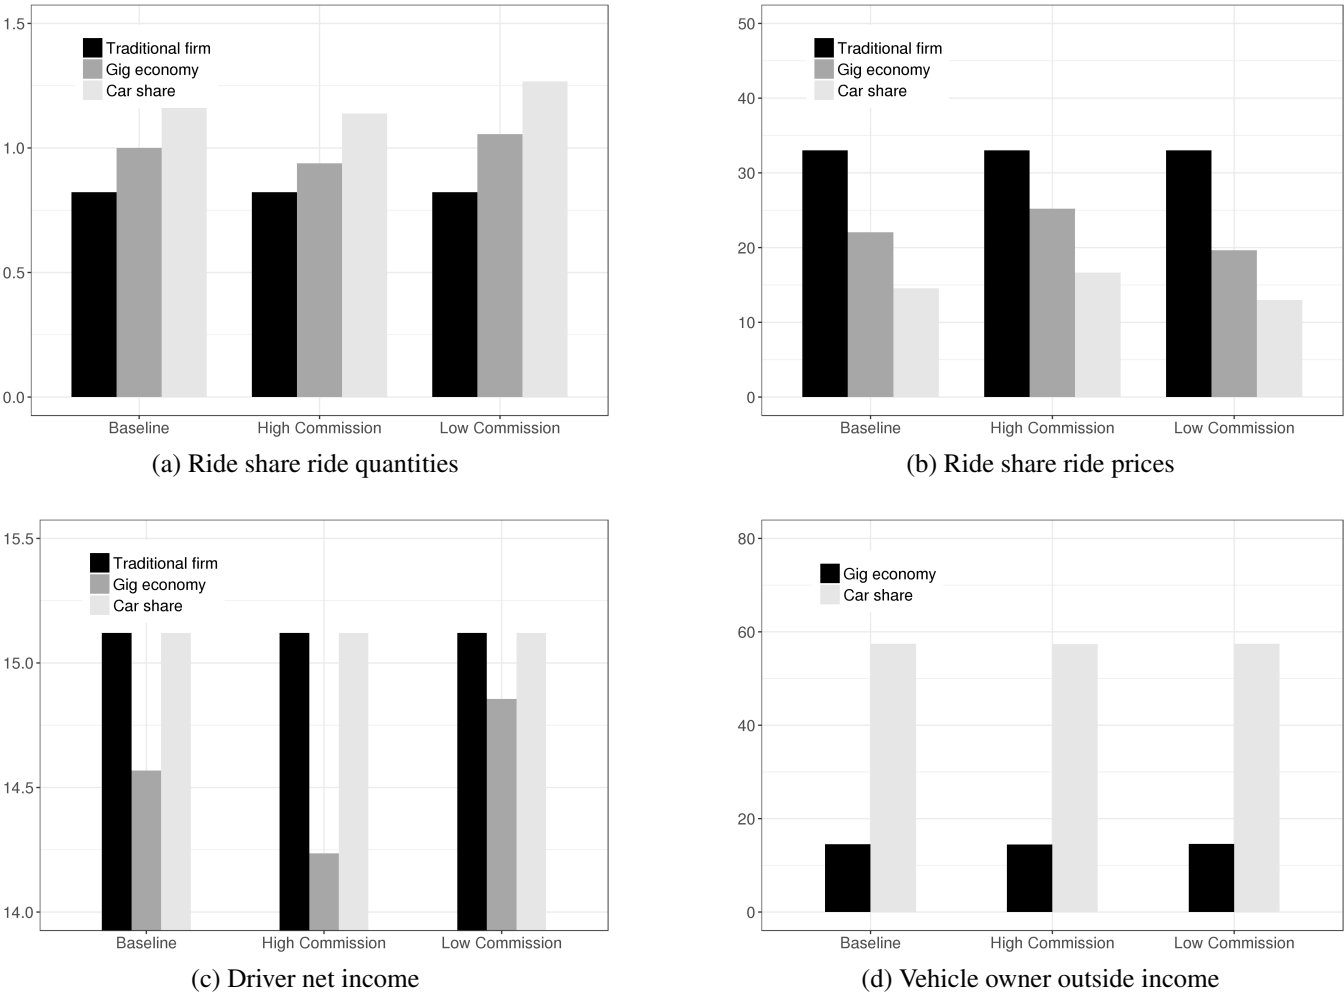


Figure C.5 compares systems of capital allocation for low (15%), baseline (25%), and high (35%) ride share commissions. The traditional firm scenario envisions firms owning financed cars and hiring drivers at minimum wage. The gig economy scenario is the current structure, where drivers work as independent 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 outside income of ride share vehicle owners.

C.7 Placebo tests

Table C.7: Nationwide auto sales and loans after ride share entry, placebo test

Panel A: New auto sales, nationwide, RL Polk data			
	<i>Dependent variable:</i>		
	Log sales		
	(1)	(2)	(3)
Post	0.0004 (0.002)	0.001 (0.002)	0.001 (0.003)
Post \times Low income	-	-0.002 (0.004)	-
Post \times High transport share	-	-	-0.001 (0.004)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr \times Low income FE	N	Y	N
Qtr \times High transport FE	N	N	Y
Observations	299,332	299,332	299,332
R ²	0.967	0.967	0.967
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Panel B: New auto loan originations, nationwide, Equifax data			
	<i>Dependent variable:</i>		
	Log new originations		
	(1)	(2)	(3)
Post	-0.002 (0.002)	-0.001 (0.002)	-0.004 (0.002)
Post \times Low income	-	-0.002 (0.003)	-
Post \times Transport share	-	-	0.003 (0.003)
Zip FE	Y	Y	Y
Qtr FE	Y	N	N
Qtr \times Low income FE	N	Y	N
Qtr \times High transport FE	N	N	Y
Observations	299,332	299,332	299,332
R ²	0.974	0.974	0.974
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table C.7 shows the placebo test corresponding to Table B.1, where entry location is randomized. Panel A is the nationwide RL Polk data; Panel B is the nationwide Equifax data. See Table B.1 for detailed specification and variable descriptions.

Table C.8: DMV registrations and liens, placebo test

Panel A: New auto registrations, select states, DMV data				
	<i>Dependent variable:</i>			
	Log new registrations			
	(1)	(2)	(3)	(4)
Post	−0.002 (0.003)	0.003 (0.003)	0.0004 (0.003)	- -
Post × Eligible	-	-	0.004 (0.003)	0.003 (0.003)
Post × Low income	-	−0.010** (0.004)	−0.008** (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
<i>Note:</i> *p<0.1; **p<0.05; ***p<0.01				
Panel B: Vehicle liens, select states, DMV data				
	<i>Dependent variable:</i>			
	Has lien			
	(1)	(2)	(3)	(4)
Post	−0.002*** (0.0005)	−0.003*** (0.001)	−0.002 (0.001)	- -
Post × Eligible	-	-	−0.002 (0.002)	0.001 (0.001)
Post × Low income	-	−0.002* (0.001)	0.001 (0.002)	- -
Post × Low income × eligible	-	-	−0.003 (0.002)	−0.005*** (0.002)
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	23,897,450	17,724,938	17,724,938	17,724,938
R ²	0.153	0.162	0.073	0.165
<i>Note:</i> *p<0.1; **p<0.05; ***p<0.01				

Table C.8 shows the placebo test corresponding to Table B.1, where entry location is randomized. Panel A is the nationwide RL Polk data; Panel B is the select state results using DMV data. See Table B.1 for detailed specification and variable descriptions.

Table C.9: Employment after ride share entry, placebo test

	<i>Dependent variable:</i>		
	log(n.filings)		
	(1)	(2)	(3)
Post	0.0002 (0.001)	−0.0001 (0.002)	(0.000)
Post × (AGI<25k)	- -	0.001 (0.002)	0.001 (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
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table C.9 shows the placebo analysis corresponding to Table (B.2) 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.

Table C.10: Vehicle utilization after ride share entry, placebo test

	<i>Dependent variable:</i>					
	Miles / year (k)					
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.127 (0.128)	0.198 (0.313)	0.461 (0.321)	- -	0.739 (0.745)	- -
Post × Low income	- -	-0.077 (0.315)	- -	- -	-0.308 (0.754)	- -
Post × Eligible	- -	- -	-0.381 (0.344)	-0.256 (0.394)	-0.629 (0.782)	-0.801 (0.833)
Post × Low income × Eligible	- -	- -	- -	- -	0.276 (0.791)	0.598 (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 C.10 shows the placebo test corresponding to Table B.3, where entry location is randomized. See Table B.3 for detailed specification and variable descriptions.

Table C.11: Zip-level credit constraints, placebo test

	<i>Dependent variable:</i>					
	log new originations		log sales		log tax filings	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	−0.002 (0.002)	−0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	−0.0002 (0.001)	−0.0004 (0.001)
Post × High bank share	0.005 (0.008)	- -	−0.004 (0.009)	- -	−0.001 (0.002)	- -
Post × High default share	- -	0.014* (0.008)	- -	0.001 (0.009)	- -	−0.0003 (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

Table C.11 shows the placebo test corresponding to Table B.5, where entry location is randomized. See Table B.5 Panels A and B for detailed specification and variable descriptions.

BIBLIOGRAPHY

- Agarwal, S., Rosen, R. J., and Yao, V. (2015). Why do borrowers make mortgage refinancing mistakes? *Management Science*, 62(12):3494–3509.
- Barrios, J. M., Hochberg, Y. V., and Hanyi Yi, L. (2018). The cost of convenience: Ridesharing and traffic fatalities. *Working Paper*.
- Benetton, M. (2017). Leverage regulation and market structure: An empirical model of the uk mortgage market.
- 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. (2012). Mortgage market design. *Review of finance*, 17(1):1–33.
- 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*.
- Cramer, J. and Krueger, A. B. (2016). Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–82.
- Dobbie, W., Goldsmith-Pinkham, P., Mahoney, N., and Song, J. (2016). Bad credit, no problem? credit and labor market consequences of bad credit reports.
- 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.

- Ganong, P. and Noel, P. (2018). Liquidity vs. wealth in household debt obligations: Evidence from housing policy in the great recession.
- Gennaioli, N., Shleifer, A., and Vishny, R. (2015). Money doctors. *The Journal of Finance*, 70(1):91–114.
- Gong, J., Greenwood, B. N., and Song, Y. (2017). Uber might buy me a mercedes benz: An empirical investigation of the sharing economy and durable goods purchase.
- Greenwood, R. and Scharfstein, D. (2013). The growth of finance. *Journal of Economic Perspectives*, 27(2):3–28.
- 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.
- 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.
- Jørring, A. T. (2017). The costs of financial mistakes: Evidence from us consumers.
- 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.
- Khwaja, A. I. and Mian, A. (2005). Do lenders favor politically connected firms? rent provision in an emerging financial market. *The Quarterly Journal of Economics*, 120(4):1371–1411.
- King, R. G. and Levine, R. (1993). Finance, entrepreneurship and growth. *Journal of Monetary economics*, 32(3):513–542.
- 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. and Sufi, A. (2012a). The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. *The Quarterly journal of economics*, 127(3):1107–1142.

- Mian, A. R. and Sufi, A. (2012b). What explains high unemployment? the aggregate demand channel.
- 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 Tchisty, A. (2010). Optimal mortgage design. *The Review of Financial Studies*, 23(8):3098–3140.
- Porta, R. L., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (1998). Law and finance. *Journal of political economy*, 106(6):1113–1155.
- Rajan, R. G. and Zingales, L. (1996). Financial dependence and growth.
- 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.