

Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform

M. Keith Chen* and Michael Sheldon

December 11, 2015

Abstract

This paper studies how the dynamic pricing of tasks in the “gig” economy influences the supply of labor. A large economic literature has explored labor supply when workers can flexibly choose how long to work each day. In a study of taxi drivers, Camerer et al. (1997) claim that drivers quit when they hit a daily income target, consequently driving less when hourly earnings are high. In general, this behavior would undermine the benefits of emerging “sharing economy” markets where tasks are dynamically priced. In this paper, we study how driver-partners on the Uber platform respond to the dynamic pricing of trips, known as “surge” pricing. In contrast to income-target findings, we find that Uber partners drive more at times when earnings are high, and flexibly adjust to drive more at high surge times. A discontinuity design confirms that these effects are causal, and that surge pricing significantly increases the supply of rides on the Uber system. We discuss the implications of these findings for earnings, flexible work, and the efficiency of dynamically-priced labor markets.

*Chen: UCLA Anderson, keith.chen@anderson.ucla.edu. Sheldon: University of Chicago, msx@uchicago.edu. Comments are welcome at keith.chen@anderson.ucla.edu. We thank Franziska Bell, Judith Chevalier, Jonathan Hall, Barry Nalebuff, Connan Snider, Molly Spaeth, and seminar participants at Princeton, RAND, UCLA, and USC for generous feedback and suggestions. All errors in this draft are the responsibility of the authors. Disclosures: Chen currently serves as the Head of Economic Research for Uber, where he designed the most recent surge algorithm. Sheldon held a summer internship at Uber. Keywords: labor supply, income targeting, dynamic pricing, ridesharing, two-sided markets, the sharing economy.

In many markets, new technologies allow traditional jobs to be divided into discrete tasks that are widely distributed across workers and dynamically priced given prevailing supply and demand conditions. This “sharing” or “gig” economy represents a shift away from fixed employment contracts to a more flexible work system, and is most common in two-sided markets in which a firm acts as a platform to connect service providers and consumers. One prominent example of this is the ride-sharing company Uber, which connects riders and driver-partners,¹ and dynamically prices trips using a system known as “surge” pricing. While there is broad consensus among economists that platforms like Uber increases consumer welfare,² less attention has been paid to its effects on the labor market in terms of earnings, work flexibility, and overall efficiency.

In the case of Uber, qualified driver-partners earn fares by providing rides on the Uber platform, and can use the platform as much or as little as they want. Drivers on the Uber platform enjoy hours flexibility, and must decide both how much and when they wish to drive. Given this flexibility, a central question is the extent to which firms can influence the supply of services on their platforms, particularly in the short term. Indeed, the income-targeting literature would suggest that perversely, increasing the price of services might actually *decrease* the supply of services on these platforms.

This paper aims to measure how the dynamic pricing of tasks on online “sharing” and “gig” economy platforms influence the supply of labor on the intensive margin.³ Central to this question are the characteristics of workers’ labor supply elasticities; that is, the responsiveness of supply hours to changes in the prevailing price of services. We measure and characterize these elasticities for drivers on the UberX platform, and find significantly and substantially positive supply elasticities. In addition, we find that increases in the “surge” price of Uber trips significantly decreases the instantaneous stopping rate of drivers on the Uber platform. That is, drivers appear to dynamically adjust their schedules to drive longer and provide more trips at times with high surge prices. This suggests that surge pricing significantly increases the number of trips that occur, and boosts the overall efficiency of the Uber system.

1 Supply Elasticities and Flexible Work

There is a large empirical literature that studies the magnitude and direction of labor supply elasticities. A substantial portion of the literature has found evi-

¹As used in this paper, a ‘driver-partner’ is someone who earns income providing rides on the Uber platform.

²Surveys of prominent economists show broad agreement that platforms such as Uber benefit consumers: http://www.igmchicago.org/igm-economic-experts-panel/poll-results?SurveyID=SV_eyDrhnya7vAPrX7,

and that in particular, surge pricing accounts for some of these benefits:

http://www.igmchicago.org/igm-economic-experts-panel/poll-results?SurveyID=SV_2bC0xviyL9hsNLL.

³Here we define the intensive labor margin is how much work a person decides to supply conditional on deciding to work that day.

dence of positive supply elasticities, in line with standard theories of intertemporal substitution of labor. These studies span a variety of markets including construction work on the Trans-Alaskan pipeline (Carrington, 1996), stadium vendors (Ottenginer, 1999), and lobster trappers/divers (Stafford, 2013).⁴ In sharp contrast to these studies however, Camerer, Babcock, Loewenstein, and Thaler (1997) finds surprising evidence of *negative* earning elasticities in a study of New York City cab drivers. Positing that taxi drivers choose their hours using reference points, Camerer et al. justified these findings theoretically by developing a theory of “income targeting”; the idea that a taxi driver has a daily income target, after which they are much more likely to stop providing rides. This results in drivers choosing to work more hours when pay is low, resulting in negative earnings elasticities. Though Camerer et al. has been an extremely influential study, it has been controversial. Several subsequent papers have found support for an income-targeting hypothesis, (Fehr and Goette 2002, Crawford and Meng 2011, Chang and Gross 2014), while others have failed to find supporting evidence (Ottenginer 1999, Stafford 2013), and others have suggested that income-target findings are largely econometric artifacts (Farber 2005, 2008, 2014).

Naturally, the direction of labor supply elasticities holds strong implications for the effectiveness of a dynamic labor-pricing mechanism. Firms such as Uber promote their dynamic pricing systems by claiming that higher fares not only temporarily lower demand but also raise supply by incentivizing drivers currently off the platform to log-on and drivers already on the system to stay on longer. To the extent that Uber partners systematically income target, the effect of higher prices on aggregate supply may be muted or even negative. That is, surge pricing could push more supply off the system than on, potentially exacerbating the supply/demand imbalance. Therefore, determining the direction as well as the magnitude of high-frequency labor supply elasticities is fundamental to the economic justification of not just Uber’s dynamic pricing system, but any firm employing a dynamic pricing system to incentivize the supply of services. Here, we study this question using data on Uber driver-partners and their response to surge pricing.

2 The Uber Marketplace and Surge Pricing

2.1 Overview of the Uber Platform

Uber is a technology firm most well-known for managing a ride-sharing platform.⁵ Uber provides a mobile application which creates a two-sided market for on-demand transportation, primarily in metropolitan areas. Riders pay a fare based on the distance of their trip and the time taken to complete the trip;

⁴For a detailed summary of the literature, see Farber (2008).

⁵While Uber offers products outside of ridesharing, for example UberEATS (on-demand meal delivery) and UberRush (on-demand courier service), we do not explore those products here.

drivers receive this fare minus a service fee paid to Uber. Payments are remitted to drivers on a weekly basis, though the amount earned is known by the driver at the end of each trip.

In this study we focus on UberX, which in the United States is Uber’s peer-to-peer service. We study UberX not only because it is the most popular, but because characteristics of other products make them less amenable to study. For example, drivers on UberBlack are licensed commercial drivers, some of who work for limo companies and may be paid a fixed salary that does not respond to surge pricing. Most have outside options which may covary with Uber market conditions,⁶ making the interpretation of their time on and off the Uber platform harder to interpret. The same criticism applies to uberTaxi drivers (taxi drivers who opt-in to the Uber platform), with the additional complication that they can also pick up street hails and are not subject to surge pricing.

Because of these issues, we focus on driver-partners who supply rides exclusively on the UberX platform.⁷ As a platform, UberX possesses many qualities which make it an ideal setting in which to analyze intensive supply elasticities. Scheduling flexibility is a feature of driving on the Uber platform; unlike most workers in the U.S. economy, UberX partners face no explicit constraints on when they can work. Conditional on qualifying to drive, an UberX partner can freely choose the days and hours they are active on the platform.⁸ Furthermore, in contrast to the taxi industry, short term vehicle leases are not typical; most UberX partners own and control their own vehicles. These characteristics avoid much of the “constrained hours” problems with measuring labor-supply elasticities in more traditional labor markets.

2.2 The Surge-Pricing Mechanism

The Uber platform adjusts its prices using a realtime dynamic algorithm known as “Surge” pricing, which has generated considerable interest among both the press and academics.⁹ Surge pricing is the output of an algorithm which automatically raises the price of a trip when demand outstrips supply within a fixed geographic area. Trip prices are adjusted by multiplying the prices of the underlying components which make up fare –the base fare, the price per mile, and the price per minute¹⁰ by a multiplier output by the surge algorithm. This multiplier is communicated to both riders and partners before each trip is initiated; riders see and must confirm, the surge multiplier (SM) before requesting

⁶UberBlack drivers may display muted earnings elasticities if low-demand conditions (low hourly earnings on Uber) correlate to low rates of commercial bookings and vice versa.

⁷Drivers can be affiliated with multiple products; for example, both UberX and UberBlack. We exclude such drivers from the analysis.

⁸For an excellent summary of what we know about the characteristics of Uber driver-partners, see Hall and Kruger 2015.

⁹See Chen et al. for an early analysis of surge pricing.

¹⁰There are some components of a fare which are fixed and not affected by the surge multiplier, such as the “Safe Rides Fee”. These fixed fees do not increase with surge, and are not remitted to the partner.

a pick-up. Partners are informed of the current surge multiplier in an area both when they are offered a pickup and also through heat-maps displayed on the partner’s mobile application (See Figure 3 below).

Surge multipliers are discrete; they range from a minimum of 1.2 up to a city-specific maximum in increment of 0.1. The specific multiplier is determined at any given time by what will be referred to as a *generator* surge multiplier (GSM), a continuous value from which the implemented surge multiplier is derived from, typically by rounding to the nearest tenth. One can think of the GSM as Uber’s estimate of the “true” market price for rides at a point in space and time, while the SM is the implemented price. There are a several rules applied to the generation of the SM from the GSM, such as maximum step sizes (how much the SM can increase or decrease from its last value) or absolute caps (a city may place a cap of 2.9 on SMs, for example). In general, unless one of these special circumstances applies the SM is the GSM, rounded to the nearest tenth. This distinction is not important for the majority of the analysis, but does play an important role in Section 4.1 when we use the rounding of GSMs into SMs as an source of identification.

3 Data and Methods

Our data represent a randomly-drawn subset of UberX partners in Chicago, Washington DC, Miami, San Diego, and Seattle. For these partners, we observe ever trip they provided on the Uber platform between September 4th, 2014, and July 4th, 2015. This comprises roughly 25 million trips.

3.1 Categorization of Sessions

A unique challenge of this dataset is the conceptualization of a partner’s “work-day” or “shift”. In contrast to papers studying the labor supply of taxi drivers, the lack of organizational constraints on time worked implies that the supply choices of Uber partners need not resemble a typical workday. While this flexibility provides driver-partners with a clear benefit, it poses an analytic challenge in defining the correct unit of time for analysis, particularly for an income-targeting hypothesis.

An intuitive first approach would be to study the decision as to how many hours a driver-partner decides to supply per calendar day. However, this approach splits in two any driving session which crosses over midnight—which, particularly on weekends, are times at which many Uber partners choose to drive. A modified but similar approach is used by Farber (2014) in his analysis of New York City taxi drivers; instead of separating a day at midnight he defines a day transition at 4AM, the time of day with the lowest number of drivers on the road.

Here, we analyze Uber partner choices using a flexible session-based approach which ameliorates many of these potential biases. We define a session as the cluster of all trip and application activity that occur without a break of more

than 4 hours. That is, a period of partner inactivity greater than four hours marks the beginning of a new session in the data. Conceptually, the use of a four-hour gap allows partners to take short breaks for meals or errands without counting such breaks as ending a session. The “length” of a session is defined as the total time that the driver is on-app in a session, either serving a ride request, or online and available to accept a trip dispatch. Given these definitions we then examine the determinants of both session length and the decision to end a session.¹¹

Figures 1 and 2 below show the average length of session per driver-partner, and how much each partner’s sessions deviate from their average session length.¹² Uber driver-partners tend to drive multiple short rather than fewer long sessions, with most sessions ranging between 2 and 5 hours, and the median driver averaging 3.47 hours per session. As figure 2 shows, session lengths also vary widely within driver. In our sample, roughly 5% of a partner’s sessions are more than twice as long as their average session, and over 18% are less than half their average. This suggests that Uber driver-partners regularly take advantage of the platform’s work flexibility. We now examine the degree to which partners use this flexibility to respond to both predictable and unpredictable changes in the price of trips.

4 Results

4.1 Autocorrelation of Income

As Farber (2005) notes, current income is only a rational input to supply decisions if it significantly predicts (immediate) future earnings. That is, your current income rate should only influence your decision to keep driving if it meaningfully predicts expected earnings going forward. To verify this interpretation of the effect of current income on supply decisions within our models, we examine the inter-temporal properties of earnings by calculating autocorrelations of the hourly city earnings.¹³ These are summarized in Table 1.

Average hourly earnings are highly correlated across hours within city in our dataset. Though the degree of autocorrelation varies by city, in all cases it is significant and substantially positive. These findings suggest that current earnings are informative of future opportunities, and that partners can use current earnings as a proxy for future earnings over reasonable time frames.

¹¹In a supplementary analysis we examine the sensitivity of our results to this definition of a “session”, and find our results are both qualitatively and quantitatively robust. Under this definition, our data are comprised of roughly 2.4 million sessions.

¹²Session length is measured as the amount of on-app time within a session. So for example, if during a session from 1 to 6pm a partner takes a 1 hour break and logs off the app, that session would be measured as 4 hours long.

¹³Here, “hourly city earnings” are the average hourly earnings for partners active in the city.

Table 1: Autocorrelations of Hourly Income

Hour	City 1	City 2	City 3	City 4	City 5
0	1.000	1.000	1.000	1.000	1.000
1	0.750	0.611	0.463	0.678	0.745
2	0.538	0.483	0.352	0.480	0.588
3	0.399	0.523	0.300	0.400	0.495
4	0.347	0.473	0.272	0.363	0.459
5	0.336	0.403	0.256	0.308	0.448

All reported statistics are significant at the 0.001 level.

4.2 The Length of Sessions: OLS and 2SLS

If partners prefer to drive when their earnings are high, then we would expect to see longer sessions at those times when realized earnings are high. Our first set of regressions estimate a simple supply hours model, where we examine the length of sessions driven by Uber driver-partners at different hourly earnings levels. The first columns of Table 2 summarize regressions which estimate the equation:

$$\begin{aligned} \log(\text{HoursOnShift}_{it}) = & \beta_0 + \beta_1 \log(\text{HourlyFares}_{it}) \\ & + \beta_2(D_i) + \beta_3(T_t) + \beta_4 \text{Weather}_{it} + \varepsilon_{it} \end{aligned} \quad (1)$$

In equation 1 the dependent variable is $\log(\text{HoursOnShift})$, where each observation is the length of an observed session. HourlyFares_{it} is calculated as the ratio of total fares earned in a session to HoursOnShift_{it} , effectively an hourly rate. The coefficient β_1 on $\log(\text{HourlyFares})$ (of which the driver-partner collects approximately 80%) can be interpreted as a measure of a driver-partners short-run supply elasticity, subject of course to a host of estimation concerns. D_i is a set of partner fixed effects, and T_t is a set of time fixed effects, including month and day-of-week fixed-effects. Finally, we include several weather controls, including temperature as a quadratic and precipitation in inches.

Common in this literature is “division bias”, the problem that HourlyFares_{it} is computed as the total fares earned by a partner in a session, divided by the number of hours on session, HoursOnShift_{it} . Because of this, any measurement error in HoursOnShift_{it} will bias our estimates of β_1 in equation 1 downward. While the Uber platform measures on-app time with extreme accuracy (up to network latency), this may still be an estimation concern if on-app time is still subject to shocks which are not a function of hourly earnings. To address this, we also estimate a two-stage model where we instrument for a partner’s hourly fares with the average hourly fares of all partners in the same city (and over the same hours), as specified in equations 2 and 3, with first-stage equation:

$$\log(\text{HourlyFares}_{it}) = \pi_0 + \pi_1 \log(\text{HourlyFares}_{-it}) + v_{it} \quad (2)$$

Table 2: The Determinants of Session Length

Variables:	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Log Own Income	0.145*** (0.00144)	0.197*** (0.00258)	0.168*** (0.00261)			
Log Average Income				0.292*** (0.00237)	0.585*** (0.00518)	0.503*** (0.00565)
Constant	1.194*** (0.00156)	1.244*** (0.00251)	1.341*** (0.00628)	1.338*** (0.00237)	1.622*** (0.00504)	1.671*** (0.00783)
Weather Controls			X			X
Fixed Effects:						
Partner		X	X		X	X
Month and Day of Week			X			X
Observations	2,377,210	2,377,210	2,368,340	2,377,210	2,377,210	2,368,340
Number of Partners	63,830	63,830	63,830	63,830	63,830	63,830
R^2	0.007	0.013	0.038			

All regressions are OLS and 2SLS regressions with Log(Session Length) as the dependant variable.

We report robust standard errors in parentheses; *** p<0.01, ** p<0.05.

and second-stage equation:

$$\begin{aligned} \log(\text{HoursOnShift}_{it}) = & \beta_0 + \beta_1 \log(\text{HourlyFares}_{it}) \\ & + \beta_2(D_i) + \beta_3(T_t) + \beta_4 \text{Weather}_{it} + \varepsilon_{it} \end{aligned} \quad (3)$$

The results for both the OLS method (Specifications 1-3) and the 2SLS method (Specifications 4-6) are given in Table 2. Specification 1 includes no controls and provides an elasticity significantly and substantially greater than zero, approximately 0.15. Controlling for partner fixed effects in specification 2 increases this elasticity slightly, which decreases again in specification 3 when calendar effects (month of year and day of week, separately) and weather (temperature in quadratic terms and precipitation totals) are added, resulting in an elasticity of about 0.17. The results from the fixed-effects (omitted here) indicate that differences between partners are significant, that partners provide more hours on the weekend, and drive less hours in response to rain or extreme temperatures.

Given the “division bias” we discuss above, we rerun the specifications in 1-3 as a two-stage model in 4-6 using the average hourly income of all partners in the same city (and over the same hours) as an instrument for a partner’s own hourly income. The specifications progress in levels of fixed effects in the same manner

as the OLS estimates. Compared to the OLS estimates, the 2SLS estimates are substantially higher; more than doubling the corresponding OLS estimates in all cases and resulting in a supply elasticity of approximately 0.50. This increase is consistent with the concern that OLS estimates may be downwardly biased due to measurement error in $HoursOnShift_{it}$.

It should be noted that the effect of instruments on our estimates, while substantial, is considerably less than in Camerer (1997) or Farber (2005). This is likely because Uber data is measured by a smartphone and displays less measurement error in active-session time. These elasticities are consistently both statistically and economically positive, so do not support the income-targeting hypothesis proposed in Camerer et al. (1997). We find that partners work longer hours when the earnings are high, and that higher prices stimulate supply on the intensive margin.

4.3 Ending a Session: Modeling the Decision to Stop

Our first analysis closely follows the existing literature, but like those papers suffers from an endogeneity problem that is difficult to overcome within a work-hours framework. Earnings levels in the Uber environment are driven by shocks to both demand and supply; and when earnings variation is driven by supply, it is likely that both outside options and/or opportunity costs for partners are strongly correlated with earnings. Therefore β_1 can not be unambiguously interpreted as an income elasticity.

Our second analysis focuses on estimating the effect of unexpected shocks to earnings, as driven by the Surge Multiplier, on a partner's decision to continue or stop driving within a session. To study this decision, we study the predictors of a trip ending a session. We take the Surge Multiplier that is in effect in the location where a partner ends a trip, as the best proxy for the earnings they should expect should they continue driving. Additional controls include cumulative measures of their session (fare, time, distance traveled, and number of trips), and current weather conditions. Conceptually, we are gauging how these factors affect a partner's decision to continue driving after each trip.

Table 3 summarize regressions which estimate the equation:

$$\Pr(EndSession_{it}) = \frac{\exp(z_{it})}{1 + \exp(z_{it})}, \quad (4)$$

where:

$$z_{it} = \beta_0 + \beta_1 SurgeMultiplier_{it} + \beta_2 CumTrips_{it} + \beta_3 CumFares_{it} + \beta_4 CumHours_{it} + \beta_5 CumDist_{it} + \beta_6 Weather_{it} + \beta_7 (D_i \times T_t \times G_{it}). \quad (5)$$

In equation 4 the dependent variable is $\Pr(EndSession_{it})$, which is the probability that an observed trip by an Uber driver-partner is their last for this session. In equation 5 the most important independent variable is $SurgeMultiplier_{it}$, which is the surge multiplier that is in effect in the geofence that a trip ends in, when that trip completes. This is the surge price that would apply to a trip that immediately begins where a partner currently is after completing their last trip. The coefficient β_1 on $SurgeMultiplier_{it}$ can be interpreted as a measure of a driver-partners short-run supply elasticity; it is the effect of multiplicative increases to trip prices on the log-odds that a partner chooses to stop driving and end a session.

Our other independent variables control for characteristics of each trip. $CumTrips_{it}$, is the cumulative number of trips already supplied by a partner in this session (and similarly for fares, hours and distance). D_i is a set of driver-partner fixed effects, and T_t is a set of time fixed effects (including Month and Day of Week fixed-effects), and G_{it} is a set of fixed effects for the geofence that a ride ends in. All fixed effects are fully-interacted. Finally, we include several weather controls, including temperature as a quadratic and precipitation in inches. Results for a Conditional (fixed-effect) Logit models are reported in Table 3, where all effects are calculated as odd ratios.¹⁴

We begin with a standard Logit in specification 1, which estimates the probability of ending a session. We see that a unit increase in the Surge Multiplier leads to a drastic decrease in stopping probability—approximately 50%. Cumulative counters in time (measured in hours) and trips have predictably significant and positive effects on stopping probabilities, as does distance (measured in miles). The effect of cumulative fare (in terms of hundreds of dollars) is significantly and substantially *positive*, indicating that partners, *ceteris paribus*, are more likely to stop when their earnings are high; a result in support of the income targeting hypothesis.

The remaining specifications are modeled as conditional Logits with increasingly restrictive fixed effects. Conceptually, by controlling for an increasingly tight set of partner, session, spatial, and temporal covariates, we can measure the effect that unexpected changes in the prevailing surge price has on the decision of Uber partners to end a driving session. In the limit these regressions

¹⁴Corresponding LPM models produce qualitatively similar results; results do not appear to be driven by changes in observations across fixed-effect levels.

Table 3: The Choice to End a Driving Session: Conditional Logits

Variables:	Conditional (Fixed-Effect) Logistic Regression			
	(1)	(2)	(3)	(4)
Surge Multiplier	0.516*** (0.00129)	0.512*** (0.00132)	0.565*** (0.00163)	0.700*** (0.00275)
C. Trips	0.955*** (0.000208)	1.033*** (0.000310)	1.008*** (0.000375)	1.039*** (0.000553)
C. Fare (hundreds)	1.248*** (0.00144)	1.091*** (0.00212)	1.000 (0.00249)	0.957*** (0.00361)
C. Time (hours)	1.176*** (0.000504)	1.156*** (0.000602)	1.110*** (0.000722)	1.119*** (0.00105)
C. Distance (miles)	1.003*** (4.02e-05)	1.010*** (5.91e-05)	1.009*** (7.28e-05)	1.003*** (0.000111)
Precipitation (inches)	0.951 (0.0292)	0.876*** (0.0287)	1.036 (0.0405)	0.797*** (0.0473)
Temperature	0.999*** (0.000175)	1.006*** (0.000208)	1.005*** (0.000240)	1.008*** (0.000330)
Temperature ²	1.000*** (1.64e-06)	1.000*** (1.96e-06)	1.000*** (2.29e-06)	1.000*** (3.22e-06)
Constant	0.124*** (0.000344)			
Fixed Effects:				
PartnerxDay		X		
PartnerxDayxHour			X	
PartnerxDayxHourxGeo				X
Observations	25,056,304	23,633,812	14,021,328	5,234,517

We report coefficients as odds-ratios and robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1.

estimate the degree to which an increase in surge prices and hence per-trip earnings, causally induces Uber driver-partners to provide more rides on the Uber platform.

Column 2 includes *Partner x Day of Week* fixed effects, controlling for the effect of a particular partner on a particular day of week. Most variables of interest remain remarkably similar—responses to surge multiplier, time, and distance are negligible. The impact of cumulative fare diminishes substantially, losing over half of its magnitude; however, it still remains significant and economically substantial in influencing a partner’s decision to stop.

Specification 3 narrows the fixed effect further to the interaction of *Partner x Day x Hour*; defining groups of observations which have multiple entries for the same partner, day of week, and hour of day. This fixed effect is quite restrictive, and as such the number of observations our model can use decreases substantially;¹⁵ nevertheless the results are still highly significant and their interpretation meaningful. Coefficients for time, distance, and trips alter slightly but not either substantially nor meaningfully. The effect of the surge multiplier decreases slightly, but still remains quite powerful. Interestingly, the effect of cumulative fare goes to 0, suggesting that the evidence in favor of the income targeting hypothesis is not robust to model specification.

Specification 4 provides an even more aggressive level of fixed effect, *Partner x Day x Hour x Geozone*. Conceptually, we are grouping observations by individual partner, day of week, hour of day, and region of city where they end the trip. That is, we are only comparing trips that the same partner ended on a Monday, between 2 and 3, in this neighborhood. Sample size decreases drastically but remains sizable. Coefficients for time, distance, and trips remain similar. The impact of surge multiplier again decreases, and this time by a much greater magnitude, but remains a powerful, negative effect on stopping probability. Cumulative fare now has a *negative* impact on stopping probability—provided the same interpretation holds on cumulative as in the previous analysis, this implies partners are less likely to stop when earnings are high. Thus, the results with respect to the impact of income on a partner’s decision to stop are consistent with the results above and neoclassical models of labor supply, and do not support the income-targeting hypothesis.

The impact the level of fixed effects have on the model highlights the powerful spatiotemporal influences that are involved with a partner’s decisions—failure to accurately capture these effects amounts to a potentially large source of bias in the estimation. The trend in coefficients displayed here tells a very compelling story—given a higher set of prices, partners will choose to work more than they otherwise would have. Thus, changes in the surge multiplier directly effect the supply decisions of partners even after they have decided to drive on a given day. Furthermore, the coefficient on cumulative fare suggests that the Surge Multiplier does not even indirectly encourage supply churn by allowing partners to hit “income targets” sooner. Rather, since partners react to higher cumulative

¹⁵The conditional Logit model requires at least one observation of each outcome within each fixed-effect group.

incomes by driving longer, surge multipliers appear to have a secondary, positive effect on labor supply through increasing incomes. Overall, it appears that the dynamic pricing mechanism is very effective in encouraging short-term supply growth on the Uber platform by encouraging partners already on the system to contribute more time than they otherwise would have.

4.4 Stopping Model with Running Variable

One concern about the results presented thus far is that, even given the strong level of fixed effects and controls already imposed, there may be underlying unobserved elements driving both the change in surge multiplier and a partner’s decision to stop; thus, the effect of the surge multiplier may be biased. In Table 4, we run the same conditional Logit models as above, but this time with a high-order polynomial in the *Generator* Surge Multiplier—the continuous number which the surge multiplier is rounded from.

By including both the SM and the GSM in the same model, the coefficient of the surge multiplier regressor can be interpreted as the average, direct impact of an increase—akin to a regression discontinuity design. That is, the high order generating variable should control *all* underlying covariates provided they do not jump discontinuously with surge prices, allowing these regressions to isolate the pure effect of the surge multiplier. The results in Table 4 suggest that surge multipliers exerts a powerful effect on the stopping probabilities of partners even under the most stringent of controls. The decrease in this effect compared to the results in Table 3 for the corresponding models can be attribute to changes in demand conditions within surge multiplier.

5 Discussion

The results presented here demonstrate the effect of increased earnings on the supply decisions of partners. In contrast to the income-targeting literature, we find that in response to surge pricing, Uber driver-partners choose to extend their sessions and provide significantly more rides on the Uber platform. This finding remains sizable even with the inclusion of extremely aggressive partner, session, spatial, and temporal controls. These controls, plus the inclusion of the generator surge multiplier, allow us to measure the supply elasticity of Uber partners in response to unexpected changes in earnings as driven by unpredictable variation in surge pricing.

Our findings suggest that Uber partners both drive at times with higher demand for rides, and dynamically extend their sessions when surge pricing raises earnings. In contrast to the existing literature, we find that Uber driver-partners do not display behavior consistent with income-targeting. These findings run contrary to a large literature on the behavior of cab drivers, which found evidence that taxi drivers reduce the supply of rides when demand is unexpectedly high. Those effects, if they held in the case of Uber surge pricing, would have significantly reduced the economic gains from dynamic pricing. To the contrary,

Table 4: Stopping Probabilities with a GSM Running Variable Control

Variables:	Conditional (Fixed-Effect) Logistic Regression		
	(1)	(2)	(3)
Surge Multiplier	0.730*** (0.00458)	0.774*** (0.00564)	0.805*** (0.00865)
GSM: 5th Order Poly.			
C. Trips	1.044*** (0.000354)	1.017*** (0.000431)	1.042*** (0.000624)
C. Fare (hundreds)	1.148*** (0.00256)	1.033*** (0.00300)	0.989** (0.00433)
C. Time (hours)	1.144*** (0.000679)	1.103*** (0.000823)	1.107*** (0.00118)
C. Dist. (miles)	1.008*** (6.91e-05)	1.007*** (8.61e-05)	1.003*** (0.000128)
Precipitation (inches)	1.068* (0.0406)	1.304*** (0.0595)	1.019 (0.0708)
Temperature	1.001*** (0.000233)	1.000* (0.000271)	1.004*** (0.000369)
Temperature ²	1.000*** (2.23e-06)	1.000*** (2.61e-06)	1.000*** (3.64e-06)
Fixed Effects:			
Partner \times Day	X		
Partner \times Day \times Hour	X		
Partner \times Day \times Hour \times Geo	X		
Observations	19,051,247	10,599,081	4,112,257

We report coefficients as odds-ratios and robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1.

we find large and pervasive positive supply elasticities, suggesting that dynamic pricing, at least in the case of Uber, significantly increases the efficiency of the ride-sharing market.

That we do not find income-targeting on the session-hours level is particularly surprising since this analysis is extremely similar to the seminal Camerer et al. and the literature that followed it. While we do not directly reanalyze their data, one possible explanation of this discrepancy is that we have extremely precise measurements of Uber partners' time on session; Farber emphasizes that measurement error in this variable has the ability to produce spurious income-targeting findings (Farber 2005, 2008). Another possibility is that on the Uber platform, earnings variation arises through extremely salient surge-induced increases in earnings-per-trip, rather than indirect earnings fluctuations through trip frequency, as in taxi markets. Finally, Uber partners interact with the platform through a smartphone interface that allows them to know current prices and session statistics like cumulative earnings, time, and trips (see Figure 3 for an example screenshot). It is possible that with access to more and more easily organized information, Uber driver-partners need not rely on rules of thumb like a daily income target.

Finally, our work is one of the first to document the degree to which platforms such as Uber enable extremely flexible work schedules, and the degree to which Uber driver-partners take advantage of that flexibility. Even under a generous 4-hour break definition, the median driver-partner averages less than three and a half hours per session, and varies that session length considerably to take advantage of surge pricing. To the degree that the "sharing economy" promises greater work flexibility, Uber driver-partners appear to take advantage of that flexibility in ways that increase their hourly earnings.

6 Conclusion

Overall, our findings support the idea that dynamic pricing significantly increases the efficiency of on-demand service markets. On the Uber platform surge pricing appears to increase the supply of rides on the Uber platform by incentivising driver-partners to provide more rides than they would have absent surge prices. We find evidence that this happens both immediately (by immediately lengthening sessions), as well as the longer time frames over which driver-partners plan their session schedules. While we investigate data on only the Uber platform, our findings suggest that dynamic pricing could significantly increase the efficiency of many of emerging "gig" markets where jobs are widely distributed across workers and in which prevailing market conditions can fluctuate across both time and location.

References

- [1] Camerer, Colin; Linda Babcock; George Loewenstein; and Richard Thaler, “Labor Supply of New York City Cabdrivers: One Day at a Time,” *Quarterly Journal of Economics* 112 (May 1997), pp. 407 – 441.
- [2] Chang, Tom and Tal Gross. “How many pears would a pear packer pack if a pear packer could pack pears at quasi-exogenously varying piece rates?” *Journal of Economic Behavior and Organization* 99 (2014) pp. 1-17.
- [3] Chen, Le; Alan Mislove; and Christo Wilson, “Peeking Beneath the Hood of Uber,” *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*, pp. 495-508.
- [4] Crawford, Vincent P. and Juanjuan Meng. “New York City Cab Drivers Labor Supply Revisited: Reference-Dependent Preference with Rational Expectations Targets for Hours and Income,” *American Economic Review* 101 (2011) pp. 1912-1932.
- [5] Hall Jonathan V. and Alan B. Krueger. “An Analysis of the Labor Market for Uber’s Driver-Partners in the United States,” *Princeton University Working Papers Industrial Relations Section*, (2015), #587.
- [6] Farber, Henry S. “Is Tomorrow Another Day? The Labor Supply of New York City Cab Drivers,” *Journal of Political Economy* 113 (February 2005), pp. 46-82.
- [7] Farber, Henry S. “Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers,” *American Economic Review* 98 (June 2008), pp. 1069-1082.
- [8] Farber, Henry S. “Why You Can’t Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers,” *Princeton University Working Papers Industrial Relations Section*, (2014), #583b.
- [9] Fehr, Ernst and Lorenz Goette. “Do Workers Work More if Wages are High? Evidence from a Randomized Field Experiment,” *American Economic Review*, 97(1) (2007) pp. 298-317.
- [10] Oettinger, Gerald S. “An Empirical Analysis of the Daily Labor Supply of Stadium Vendors,” *Journal of Political Economy* 107 (April 1999), pp. 360-392.

7 Figures

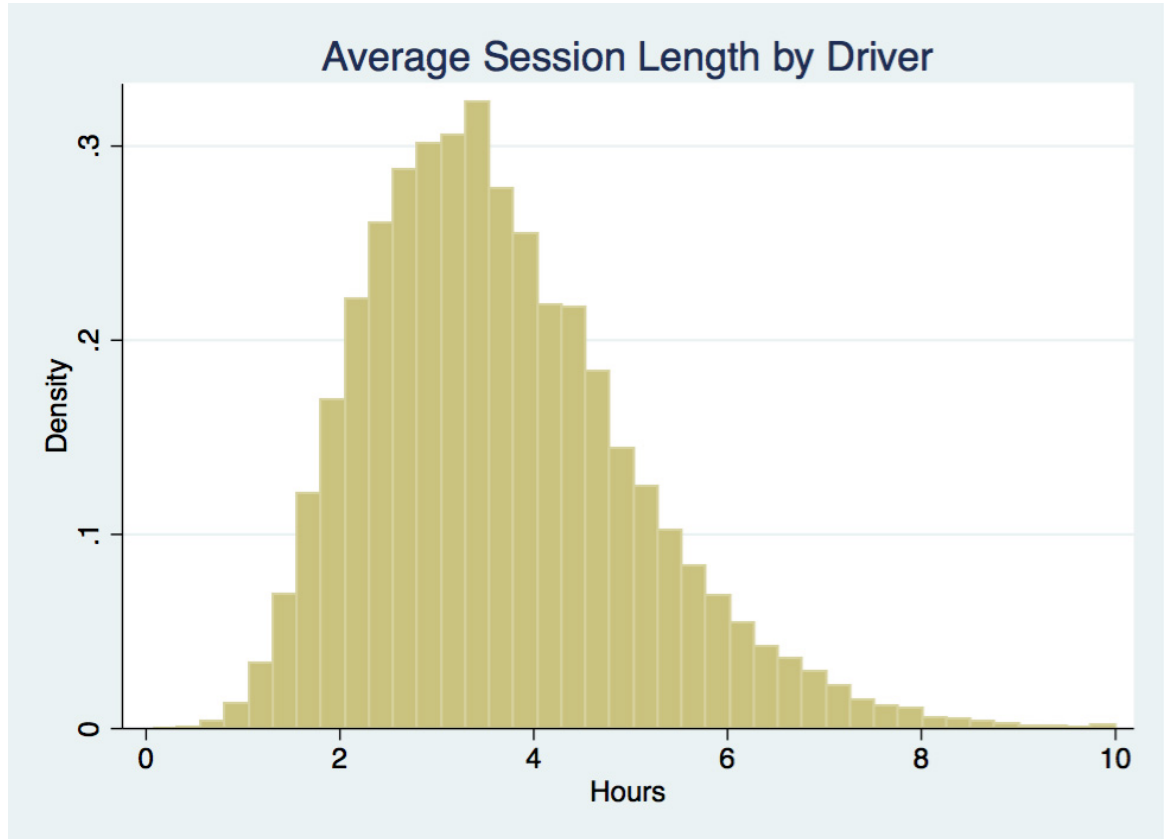


Figure 1: Uber Partners' Average Session Lengths

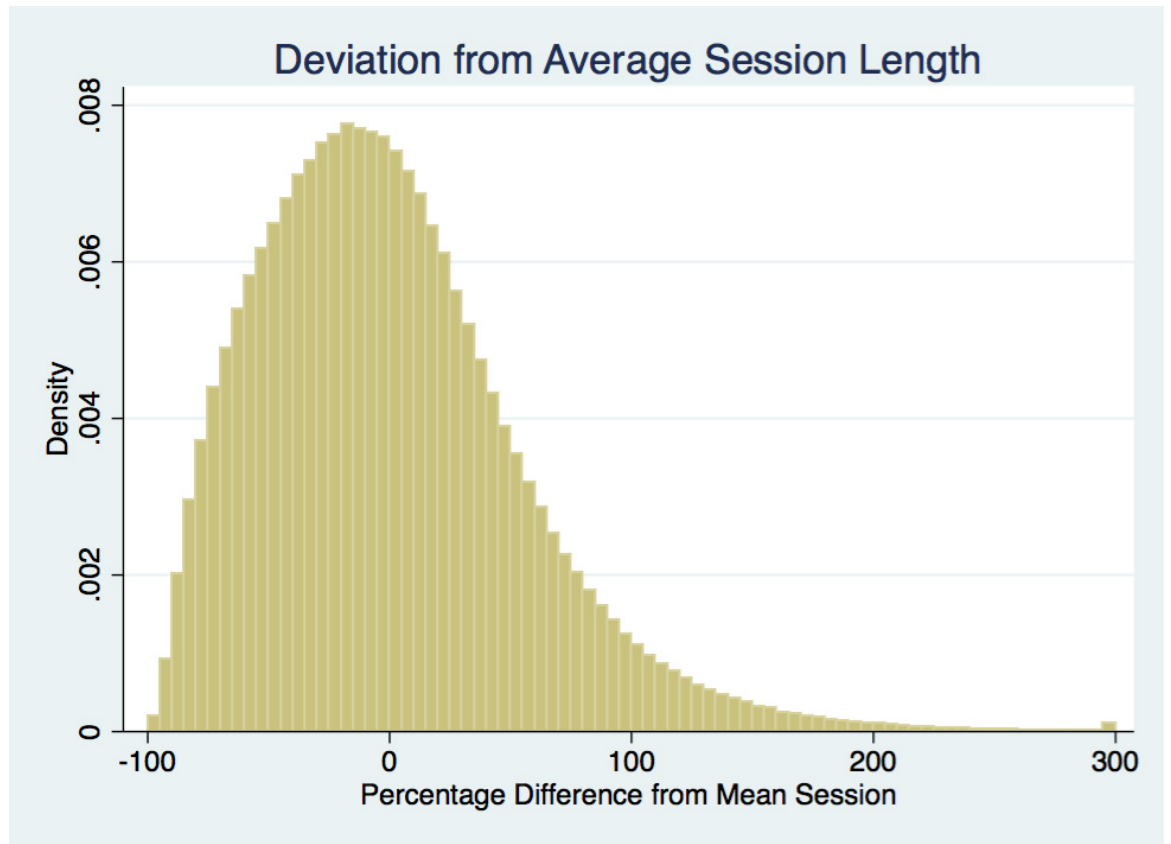


Figure 2: Variation of Session Length from a Driver's Average

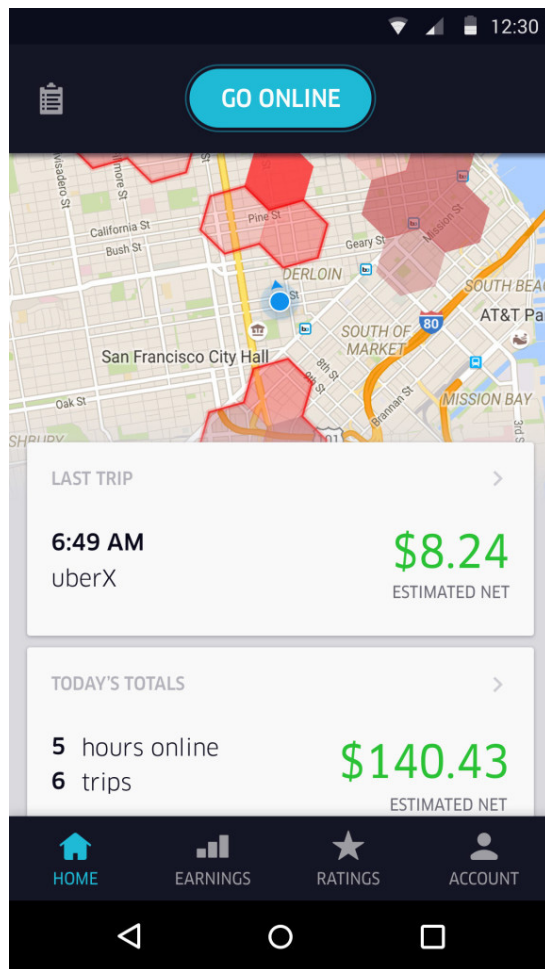


Figure 3: Tools on the Uber driver-partner smartphone application.