Do Digital Platforms Reduce Moral Hazard? The Case of Taxis and Uber *

Meng Liu †
(WUSTL and MIT)

Erik Brynjolfsson ‡
(MIT and NBER)

Jason Dowlatabadi §
(Uber Technologies)

July 15, 2018

Abstract

Digital platforms have made possible a variety of technology-enabled tools that enhance market transparency, including real-time monitoring, ratings of buyers and sellers, low-cost complaint channels, and new pricing schemes. How do these innovations affect moral hazard and service quality? We investigate this problem by comparing driver routing choices and efficiency at a large digital platform, Uber, with traditional taxis. The identification is enabled by matching taxi and Uber trips at a very granular level so they are subject to the same underlying optimal route, and by exploiting characteristics of the pricing schemes that differentially affect the incentives of taxi and Uber drivers in various circumstances. We find that (1) taxi drivers detour on airport routes by an average of 7.4%, with non-local passengers on airport routes experiencing even longer detours; (2) taxi drivers overall drive at a greater speed than Uber drivers; and (3) Uber drivers are more likely to detour during periods of high surge pricing. These findings are consistent with the platform tools reducing driver moral hazard, but not with competing explanations such as driver selection or differences in driver navigation technologies.

*We thank Keith Chen, Dean Eckles, Andrey Fradkin, Xiang Hui, Tom Lam, and John Horton for their invaluable comments and suggestions. We thank seminar participants at MIT, Uber, AEA 2018, Marketing Science 2018, and SICS 2018. The MIT Initiative on the Digital Economy provided generous research support and Uber provided essential data. Dowlatabadi is a current employee at Uber. The views expressed here are those of the authors and do not necessarily reflect those of Uber Technologies, Inc. All errors are ours.

† mengliu@mit.edu
‡ erikb@mit.edu
§ jasond@uber.com

1 Introduction

The economic effects of digital platforms are rapidly growing, especially in the service sector. Examples include Uber for ride-hailing, Airbnb for accommodations, Rover for dog-walking, Urbansitter for baby sitting, etc. Do digital platforms significantly affect moral hazard or service quality, compared to traditional settings? Answer to this question is of crucial importance for a better understanding of the nature of online-offline competition and welfare in the digital economy. In this paper, we study this question by comparing a particularly successful and pervasive digital platform, Uber, with traditional taxis.

Specifically, we investigate driver detours and travel speed, two key metrics of driver routing efficiency and service quality. Detour in our context is defined as the extra distance a driver adds to the optimal passenger route. This is a measure of driver moral hazard. In a hypothetical situation where a taxi driver and an Uber driver drive between the same two points at the same time, the difference in their routing decisions should reflect differences in their payoffs, including the benefits and costs of detours and speeding. To the extent that features such as GPS navigation, tech-aided monitoring, ratings, and digital feedback increase market transparency for passengers and therefore increase penalty of driver moral hazard, the Uber driver's routing is likely more efficient than that of the comparable taxi driver in situations with high moral hazard payoffs for both drivers.

There are two key challenges in identifying driver moral hazard. First, without directly observing driver moral hazard,[1] one needs to construct valid counterfactuals to infer opportunistic behavior by using detailed trip-level data of both taxis and Uber. We overcome this data challenge by combining taxi and UberX trip records in New York city (NYC) and matching taxi and Uber trips at granular route-time level. As a result, the drivers of the matched trips are plausibly subject to the same real-time optimal routes, even if they are not directly observed.

These matched pairs of taxi and Uber trips then become our units of analysis, where we explore the variation in the within-match taxi-Uber difference in trip distance and duration

[1] The inability to directly observe driver moral hazard is due to the lack of the optimal routing benchmark at the time of the trip. For example, using a long-run average trip distance queried from routing engines such as Google Maps may under-estimate the true real-time optimal route and over-estimate the detour if there was a road closure on the optimal route.

across route types that represent different moral hazard incentives. However, the estimation approach (essentially a "diff-in-diff" approach) would face an identification challenge if route types were not randomly assigned to taxi and Uber drivers. We address this challenge by exploiting the institutional background of the taxi industry and the Uber platform. For taxis, there is no passenger selection of taxi drivers as taxi drivers are *ex ante* homogeneous to passengers. Due to the strict taxi refusal law and sample over-representation of short trips in thick markets and airport trips, driver selection of passengers is at best limited. On Uber, rider assignment, performed by Uber's algorithm, is practically random to individual drivers, and driver selection of passengers is deterred through multiple mechanisms. Therefore, the market itself approximates the experimental ideal. Nonetheless, we also exploit the within-driver variations in some specifications to further purge potential selection issues.

We find that drivers indeed respond opportunistically to changes in incentives. When the fare is metered in trip distance, taxi drivers detour by an average distance of 7.4% on airport trips. Taxi drivers detour even more when the rider is non-local on a metered airport trip. Uber drivers are not immune to moral hazard either. We find that they are more likely to detour on metered airport trips when they are paid high surge multipliers. When we compare trip durations, we find that taxis, on average, travel faster and finish trips earlier than Uber drivers, which is consistent with a rational response to their different pricing schedules.

These findings fit a model of driver moral hazard, where the driver decides on the amount of detour and speed of travel, given the set of information, the pricing schedule, and overall incentive devices at work. The key tension is a trade-off between payoff and cost of the opportunistic behavior, where the cost includes expected monetary and reputation cost of opportunistic behavior, and the expected forgone earnings lost due to the opportunistic behavior (for example, detour usually prolongs travel time, reducing opportunities for additional trips). It then follows that drivers lack detour incentives when driving short trips in thick markets (e.g. within-Manhattan trips), in light of low returns and high opportunity costs, and the detour incentive is greater on airport trips. On the other hand, the driver's speeding incentive increases when the metered fare does not also compensate for driving time.

We strengthen our identification of driver moral hazard by exploring competing explanations. First, we demonstrate that this is not mainly driven by drivers self selecting into more profitable routes. This is done by observing no significant change in estimation results when Uber or taxi driver fixed effects are controlled for. Second, driver selection can also take place on the extensive margin. Following 1,549 former taxi drivers who switched to Uber, we find that these drivers used to detour as taxi drivers and their detour behavior was no longer observed after joining Uber. This is strong evidence that drivers re-optimize under Uber's arrangement via behavioral change. Third, we show that the data are not compatible with the hypothesis that the difference in GPS adoption accounts for the observed taxi driver routing inefficiency. Lastly, we examine whether or not the observed detouring behaviour could be the result of optimizing for travel time as opposed to distance. We rule out this possibility by demonstrating that taxi drivers arrive later than comparable Uber drivers when detouring and the extra travel time increases with the amount of detour.

This industry offers an ideal setting to study a research question like ours. First of all, it is a textbook example of competitive marketplace, where both taxi and Uber drivers offer a homogeneous, well-defined service (namely, transporting a passenger from one point to another), take the price as given, and maximize earnings as individual entrepreneurs. Therefore, full exploitation of opportunities can be safely assumed. Second, GPS coordinate data allow us to make valid comparisons between taxis and Uber at the trip level. Such counterfactual groups can be difficult to form in other industries. Lastly, identification of moral hazard is facilitated by the institutional features of taxis and Uber, where driver behavior responds to exogenous changes in route characteristics. Therefore, we make a strong case by drawing evidence from the taxi industry, and our findings can shed light on other marketplaces as well.

1.1 Literature and Contribution

Our paper is closely related to several strands of the literature. The first is on how technology, particularly information technology (IT), mitigates the agency problem in various settings (Tabarrok and Cowen (2015)). In the typical workplace, IT-enabled monitoring is found to be productivity-enhancing through complementing performance pay (Aral et al. (2012), Bresnahan et al. (2002)), reducing employee shirking (Nagin et al. (2002)) or misconduct (Pierce et al. (2015)), and increasing standard process compliance (Staats et al. (2016)). In the context of trucking, Hubbard (2000) finds that on-board computers that facilitate monitoring of drivers increase productivity by improving both driver incentives and managers' resource allocation decisions. Duflo et al. (2012) show that incentive pay enabled by tech-aided monitoring can raise teachers' attendance rate and consequently student performance. Reimers et al. (2018) find that insurance companies' monitoring technologies reduce driver moral hazard and fatal accidents. Sudhir and Talukdar (2015) illustrate the role of IT at inducing business transparency by showing more corrupt businesses tend to resist IT adoption more.

Besides the traditional settings, there are also studies on digital market designs that improve productivity by regulating agent incentives. Hui et al. (2016) identify efficiency gains from eBay's buyer protection program as a result of reduced seller moral hazard and seller adverse selection. Klein et al. (2016) show that a change in eBay's policies that lead to less biased buyer ratings of sellers also improved seller effort and quality without inducing sellers to exit the market. Gans et al. (2017) evaluate the role of Twitter as a mechanism of consumer voice in disciplining firms for low quality. Liang et al. (2016) find that IT-enabled monitoring mitigates moral hazard on an online labor platform.

While these aforementioned studies focus on technological improvements either within the offline setting or within the online setting, we are among the first to provide a direct online-offline comparison to study the relationship between technology, agent incentives, and quality provision. As many sectors are being digitized, empirical studies of how incentives and quality provision differ between online and offline markets become crucially important for a better understanding of competition and welfare in the digital economy.

Second, this paper relates to the literature on digital disruption and online-offline competition (Bakos (1997), Brown and Goolsbee (2002), Brynjolfsson et al. (2003), Brynjolfsson and Smith (2000), Forman et al. (2009), Waldfogel (2017), among many others. See Goldfarb and Tucker (2017) for a review.). In particular, this paper contributes to the studies of emerging tech-aided ride-hailing platforms. These platforms, compared to traditional taxis, enable more efficient matching of drivers and passengers with real-time technologies and dynamic pricing (Castillo et al. (2017), Guda and Subramanian (2018), Hall et al. (2015)), as reflected in higher driver utilization (Cramer and Krueger (2016)) as well as quick adjustments to market equilibrium (Hall et al. (2017)). Specifically, efficiency induced by dynamic pricing critically depends on consumer preferences and the tradeoff between wait time and price (Lam and Liu (2017)), as well as driver labor supply that responds to wage fluctuations (Chen and Sheldon (2016)). Consumers benefit from ride-hailing platforms extensively (Cohen et al. (2016)). Drivers also benefit from these platforms due to flexible work arrangement (Chen et al. (2017), Hall and Krueger (2015)) and commission schemes that allow for driving without a lease (Angrist et al. (2017)). In this paper, we find that these technological and organizational features have important implications on driver incentives and quality provision, and thus add an important layer in the analysis of efficiency.

Finally, our findings resonate with empirical work on taxi driver opportunistic behavior. Balafoutas et al. (2013) identify taxi driver detours when passengers are less informed about the optimal routes or the local taxi fare structure, via a field experiment in Athens, Greece. Also studying NYC taxi market, Liu et al. (2017) identify likely non-locals from locals based on the destinations of trips originating at New York's airports. They find that taxi drivers defraud non-locals more on LaGuardia trips that are metered, but not so on JFK flat-fare trips. Balafoutas et al. (2017) show that drivers may also defraud more when passengers explicitly state that their expenses will be reimbursed, giving rise to a second-degree moral hazard. Rajgopal and White (2015) point out the importance of regulatory restrictions on driver fraud as they find greater likelihood of driver fraud when dropping passengers off in areas where taxis are not allowed to pick up subsequent passengers. We contribute to this literature by demonstrating that driver moral hazard can be mitigated by tech-aided ride-hailing platforms.

2 NYC Taxis vs. Uber: Pricing and Market Design

2.1 Taxi and Uber Pricing

NYC taxi fares are set by Taxi and Limousine Commision (TLC)[2]. Most routes are metered with a base fare of \$2.50 upon entry, plus a linear component that is \$0.50 for every $\frac{1}{5}$ miles traveled, as well as taxes, fees, and tolls incurred. A \$0.50 per-minute charge is applied in place of the per-mile charge when the traffic is slow (under 12 miles per hour). Routes between Manhattan and JFK airport are not metered. Instead, a flat rate of \$52 plus taxes, fees, and tolls applies. Some taxi drivers are medallion-owners who essentially run the business as an entrepreneur, while other drivers pay a daily, weekly, or monthly lease to medallion owners and keep all revenues minus gas and some maintenance.

Unlike taxis, Uber does not differentiate between fast and slow traffic. The UberX base fare includes a fixed component of \$2.55 upon entry, \$0.35 per minute of travel, and \$1.75 per mile of travel, plus taxes, fees, and tolls. On top of the base fare, passengers also need to pay whatever the surge multiplier in effect at the time of request. For a 2-mile, 10-minute trip with a surge multiplier of 2, UberX costs $2 \times (\$2.55 + \$0.35 \times 10 + \$1.75 \times 2) = \$19.10$, plus taxes, fees, and tolls. Designed to balance real-time demand and supply (Hall et al. (2015)), Uber's surge multipliers change rapidly across locations and time (Lam and Liu (2017)). Unlike taxi fixed fare on certain routes, all Uber routes in NYC are metered according to the same pricing formula. Uber drivers keep all trip earnings minus commission fees, which usually run between 20% and 30%, and they are responsible for all operation-related expenses, such as insurance, maintenance, and gasoline.

2.2 Taxi and Uber Market Design

The market design for the Uber platform differs significantly from taxis. First, GPS navigation is widely adopted and used by Uber drivers, while taxi drivers mainly navigate without GPS. The Uber app is designed in a way that GPS navigation becomes an integral part: when the driver picks up a passenger and starts the trip, Uber's built-in GPS automati-

[2] Refer to the official language on the pricing rule: http://www.nyc.gov/html/tlc/html/passenger/taxicab_rate.shtml

cally initiates, or the app switches to the preferred GPS that the driver has set up (eg. GoogleMaps and Waze).

Second, Uber implements a set of market designs that aim at aligning driver incentives, and these market designs are absent or costly with taxis. With the Uber app, passengers can readily monitor driver routing in real time – passengers can either monitor the route taken through their own smart phone app, or directly seeing the driver's app since driver's phone is usually mounted in a way that it is visible to passengers. This way, passengers can easily tell whether or not the driver is taking the route that is given by the GPS.

Uber ratings are designed to be very easy for users to access (i.e. a swipe of stars and a click to confirm) and therefore most passengers do rate their drivers (73.5% for NYC UberX, January - June 2016). Similar to other reputation systems, Uber driver ratings are highly concentrated with a mean of 4.74 (see Figure 1). Drivers with low ratings are constantly warned by Uber. Uber starts to consider deactivating a driver when the driver rating is below a threshold (4.5 in NYC). Drivers appear to take ratings seriously[3] and perhaps as a result, the deactivation risk is quite low (about 3%).

Figure 1: NYC UberX Driver Rating Distribution, Jan-June 2016

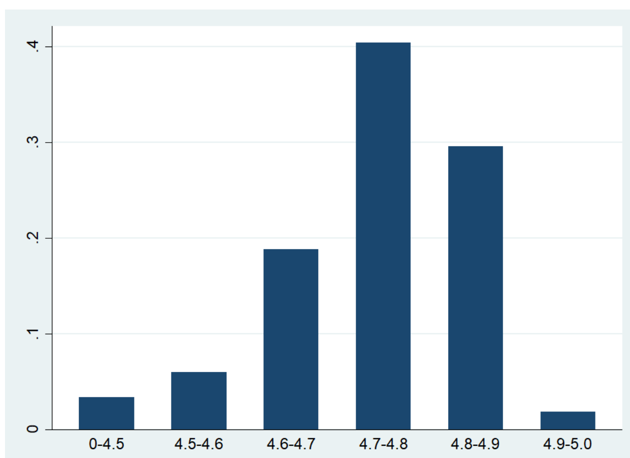

In addition to monitoring and rating, verification and complaints are also made easy on Uber due to electronic trip records stored. Passengers can re-visit the historical trip summaries in their app to check and verify certain details. In the case of negative riding

[3]The qualitative study by Lee et al. (2015) states that "Drivers took their ratings seriously. High ratings such as 4.98 became a source of pride whereas a rating below 4.7 became a source of disappointment, frustration, and fear of losing their jobs."

experience, Uber passengers can easily file a complaint through the app and Uber customer service handles the conflict resolution by evaluating the trip records. By contrast, taxi passengers in these situations can either call the TLC hot line or visit their website, but the process is usually long and may require legal procedures. In 2016, taxi complaints are 1 in every 6,356 trips, whereas Uber fare adjustments are more than 30 times more likely: 1 in every 170 trips. Figure 2 lists the main reasons of fare adjustments, with the number one reason being "inefficient route".

Uber's market design enhances market transparency and, at least in theory, should promote market efficiency. In particular, timely and effective monitoring and regulation should be able to discipline agent behavior by reducing moral hazard.

Figure 2: Uber Fare Adjustment Reasons

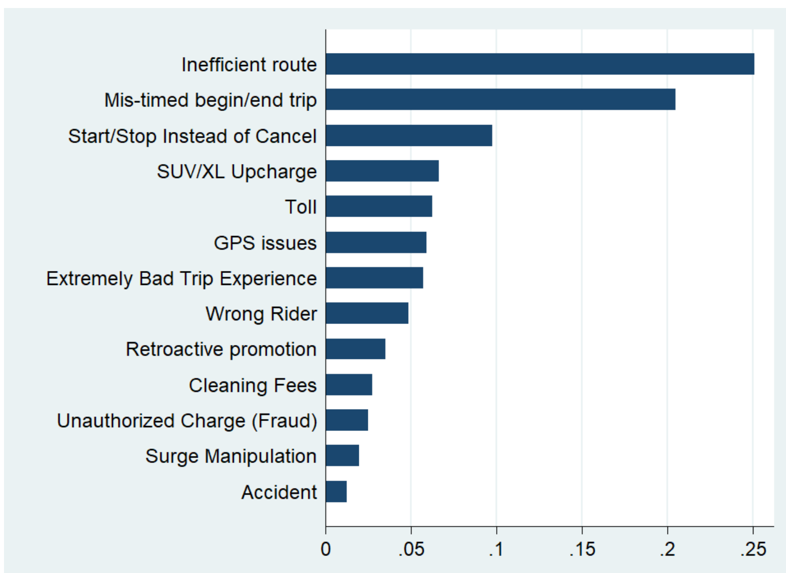

3 A Theoretical Framework of Driver Moral Hazard

In this section, we describe a theoretical framework of driver moral hazard that builds on Liu et al. (2017), where a representative, risk-neutral driver maximizes her payoff by deciding on the amount of detour and driving speed. For a given pair of pick-up and drop-off locations whose route length is $\underline{d}$ (as given by a map, say), the realized trip distance, d, is represented by the following equality:

$$d = \underline{d}(\Delta d + x + \epsilon_d), \tag{1}$$

where Δd represents the driver's navigation skill, e.g., driver's knowledge of the streets, $\Delta d \in [1, +\infty)$. Let x denote the amount of detour, where $x \in [0, +\infty)$. Let ϵ_d denote the random driver-route shock that affects routing efficiency, which is normally distributed with a mean 0 [4].

Similarly, the trip duration, t, is defined as follows:

$$t = \underline{t}(\Delta t + \lambda x + y + \epsilon_t), \tag{2}$$

where $\underline{t}$ is the expected trip duration for a route of length $\underline{d}$, based on the on-going traffic of the route. The driver's driving skill in reflected in Δt, where $\Delta t \in [1, +\infty)$. Let λx measure the additional travel time because of the detour x, where $\lambda \in [0, +\infty)$. Let y represent the extra travel time incurred when the driver drives at a speed other than implied by $\underline{t}\Delta t$: $y > 0$ when the driver drives relatively slow, and $y < 0$ when the driver drives relatively fast. Let ϵ_t denote the random driver-route shock that affects trip duration, which is normally distributed with a mean 0.

Under a linear pricing rule, the driver chooses the amount of detour (x) and the speed of driving (equivalent to y) to maximize the following expected payoff function:

$$\begin{aligned}\underset{x,y}{\text{Max}}\ \text{E}\ \{&\gamma[p_0 + p_d\underline{d}(\Delta d + x + \epsilon_d) + p_t\underline{t}(\Delta t + \lambda x + y + \epsilon_t)] \\ &- f(x; \underline{d}, \Theta_f) - g(y; \underline{t}, \Theta_g) - q_e\underline{t}(\lambda x + y)(\frac{p_0 + p_d D_e + p_t T_e}{T_e})\},\end{aligned} \tag{3}$$

where γ denotes the surge multiplier: it is equal to 1 for taxi trips and greater than or equal to 1 for Uber trips; p_0 is the base fare upon entry; p_d is the rate per unit of distance; p_t is the rate per unit of time (note that in normal traffic, $p_t = 0$ for taxis). Therefore, the first part of the maximand is the total earning of the trip.

[4] It is possible that the realized trip distance is shorter than the map distance $\underline{d}$, when the random shock ϵ is sufficiently negative. This occurs, for example, when a road turn is permitted during certain time of the day, which shortens the route but is not captured by the map distance.

When detouring, the driver incurs a cost f, which can be viewed as the probability of getting caught times the monetary and/or reputation penalty of the detour. The cost may be in the form of fines (taxis), lost tips (taxis) [5], low ratings (Uber), and refund to passengers (taxis and Uber). f is assumed twice differentiable in x, with route length $\underline{d}$ and a parameter set Θ_f, and $f(x=0)=0$, $f_x>0$, $f_{xx}>0$, as longer detour is more likely to be caught and penalized more. Moreover, it is assumed that detour is progressively costly on longer routes, or $f_{x\underline{d}}>0$, $f_{xx\underline{d}}\geq 0$. In addition, $\gamma \in \Theta_f$, and $f_{x\gamma}>0$, meaning that the marginal detour penalty on Uber is greater when surge is greater. Defined similarly as f, g is the expected monetary and/or reputation cost associated with the trip duration: $g>0$ for all y, i.e., both driving unnecessarily slow and unnecessarily fast tend to be noticed and penalized by the passenger; $g_y<0$ for $y<0$ and $g_y>0$ for $y>0$; finally, $g_{yy}>0$ and $g_{y\underline{t}}>0$.

Let q_e represent the probability of getting a subsequent passenger at the drop-off location and time if there was no detour; $\underline{t}(\lambda x+y)$ measures the forgone service minutes due to the detour; $\frac{p_0+p_dD_e+p_tT_e}{T_e}$ represents the per-minute earning of the forgone trip, where D_e and T_e are the expected length and duration of the forgone trip, respectively. Thus, the last part of the maximand represents the opportunity cost of detour in terms of the forgone payoff that could have been earned from a subsequent trip.

Taken together, the driver's problem in Equation 3 is to solve two trade-offs: one trade-off is between the monetary reward of detour and the opportunity cost of detour, which consists of the expected detour penalty and the forgone payoff because of the detour; the other similar trade-off applies to driving speed. The first-order conditions are listed below,

$$f_x(x;\underline{d},\Theta_f)+q_e\underline{t}\lambda\left(\frac{p_0+p_dD_e+p_tT_e}{T_e}\right)-\gamma(p_d\underline{d}+p_t\underline{t}\lambda)=0, \tag{4}$$

$$g_y(y;\underline{t},\Theta_g)+q_e\underline{t}\left(\frac{p_0+p_dD_e+p_tT_e}{T_e}\right)-\gamma p_t\underline{t}=0. \tag{5}$$

Then the following comparative statics follow:

(1) $\frac{\partial x^*}{\partial \underline{d}}=\frac{\gamma p_d-f_{x\underline{d}}}{f_{xx}}\lesseqgtr 0$, depending on the sign of $\gamma p_d-f_{x\underline{d}}$. That is, the driver detours more (less) on longer routes when the increase in marginal detour profitability due to longer

[5] By the end of our sample period, Uber had not implemented the tip feature in the application.

routes is greater (less) than the increase in marginal detour penalty due to longer routes. Similarly, $\frac{\partial y^*}{\partial \underline{t}} = \frac{\gamma p_t - q_e(\frac{p_0+p_d D_e+p_t T_e}{T_e}) - g_{y\underline{t}}}{g_{yy}} \lesseqgtr 0$.

(2) $\frac{\partial x^*}{\partial \theta} = -\frac{f_{x\theta}}{f_{xx}} < 0$, for a parameter $\theta \in \Theta_f$ that increases the marginal detour penalty ($f_{x\theta} > 0$). For example, the driver detours less (more) when the rider is a local (non-local) passenger who is more likely to notice the detour. Similarly, $\frac{\partial y^*}{\partial \theta} = -\frac{g_{y\theta}}{g_{yy}} < 0$, for a parameter $\theta \in \Theta_g$ that increases the marginal detour penalty ($g_{y\theta} > 0$).

(3) $\frac{\partial x^*}{\partial \gamma} = \frac{p_d \underline{d} + p_t \underline{t} \lambda - f_{x\gamma}}{f_{xx}} \lesseqgtr 0$, depending on the sign of $p_d \underline{d} + p_t \underline{t} \lambda - f_{x\gamma}$. This means that the driver detours more (less) when the surge multiplier is higher if the increase in marginal payoff of detour due to high surge is greater (less) than the increase in marginal detour penalty due to high surge. Similarly, $\frac{\partial y^*}{\partial \gamma} = \frac{p_t \underline{t} - g_{y\gamma}}{g_{yy}} \lesseqgtr 0$, depending on the sign of $p_t \underline{t} - g_{y\gamma}$: when the increase in marginal payoff of additional travel time because of surge exceeds the increase in marginal penalty of additional travel time because of surge, Uber drivers have the incentive to drive for longer time.

(4) $\frac{\partial x*}{\partial q_e} = -\frac{\underline{t}\lambda(\frac{p_0+p_d D_e}{T_e})}{f_{xx}} < 0$. Other things held constant, the driver detours more (less) when the demand at the drop-off location is lower (higher), due to a lower (higher) opportunity cost of detour. Similarly, $\frac{\partial y*}{\partial q_e} = -\frac{\underline{t}(\frac{p_0+p_d D_e}{T_e})}{g_{yy}} < 0$.

(5) In normal traffic, $y^{c*} < 0$, and $y^{c*} \leq y^{u*}$, where c denotes taxi cabs and u denotes Uber. Everything else held constant, taxi drivers have greater incentives than Uber drivers to drive faster than the ongoing traffic on the road, because they are not paid by their driving time ($p_t = 0$).[6]

4 Data and Sample Construction

4.1 Data

Our data combines NYC taxi trip records and UberX trip records for two six-month periods: January to June, 2016 and July to December, 2013. Taxi trip records contain detailed information such as pick-up and drop-off time and GPS coordinates, trip distance and duration, and fares and fees of various sorts. The 2013 taxi data contain anonymized driver ID

[6]In the rare case of slow traffic with sufficiently low drop-off demand, it is possible that $y^{c*} > 0$ and $y^{c*} \geq y^{u*}$ because $p_t > q_e(\frac{p_0+p_d D_e+p_t T_e}{T_e})$.

and medallion numbers, but these identifiers were removed later on by TLC due to privacy concerns. UberX trip records contain similar information, plus extra information such as the surge multiplier, anonymized driver ID, driver total number of trips on Uber, anonymized rider ID, rider total number of trips on Uber, driver lifetime rating, driver and rider rating of a given trip.

4.2 Matching of Comparable Taxi and Uber Trips

To make a valid comparison of taxi and Uber routing, we match taxi and Uber trips at granular route-time levels, such that the matched trips are subject to the same underlying optimal routing. In brief, we match an Uber trip to a given taxi trip if they share the same pick-up location and dropoff location, and their start times are within a short time window of each other. In the remainder of this section, we detail the steps of the matching process.

Figure 3: Matching of Taxi and Uber Trips

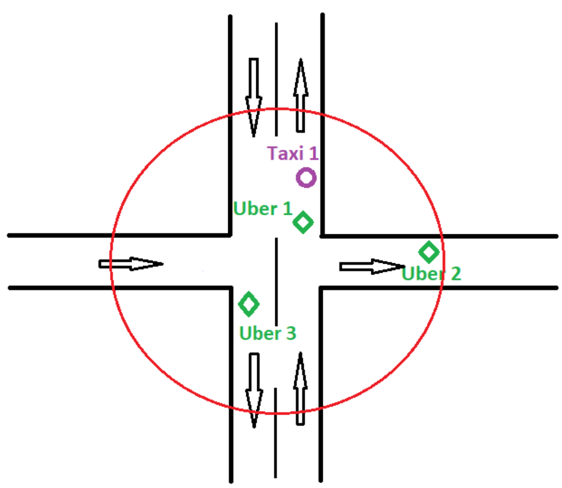

Step 1: Because of the exceedingly high concentration of pick-ups and drop-offs around street intersections, we first define locations by dividing NYC into small Voronoi cells centered at street intersections, where each street intersection is approximately 100 meters from its closest neighboring intersections. An example of these Voronoi cells is given by Figure 9. Using Figure 3 as an illustration, this means that we match Taxi 1, Uber 1, Uber 2,

and Uber 3 in the circled area. However, we show in Figure 10 that if we did nothing else, this matching criterion would yield a stark difference in the distribution of pick-ups between taxis and Uber – taxi pick-ups (purple) are more concentrated on major avenues and streets, whereas Uber pick-ups (green) are more from cross-town streets with relatively slow-moving traffic. Similar distribution applies to matched drop-offs as well. This may reflect the difference in drivers searching and matching with passengers: taxi drivers mainly cruise on major avenues and streets to look for passengers, while Uber drivers more often pick up passengers at their door steps.

Step 2: We then restrict matches to be on the same street, because trips on different streets are subject to different optimal routes, therefore they cannot be valid counterfactuals to each other. In the example of Figure 3, this means that Taxi 1 will be matched with Uber 1 and Uber 3, instead of Uber 2.

Step 3: Following a similar logic as Step 2, we further filter out matches that follow different traffic directions of the same streets. Therefore, Uber 3 ceases to be matched with Taxi 1, and Taxi 1 and Uber 1 of Figure 3 remain in the sample so far as they are in the same traffic direction.

We then apply the same filters (Step 1, 2, 3) for drop-offs as well.

Step 4: We then keep matched pairs whose pick-ups are within a short time window from each other, so that the matched trips are subject to the same traffic, road conditions, as well as other common factors. The time window for the main analysis is set at 15 minutes, and we apply various time windows (eg. 5 minutes, 10 minutes, 20 minutes) in the robustness checks.

Finally, we discover that in the raw TLC taxi trip records, there are two taxi meter vendors with about equal shares, where Vendor 1 reports trip distance to the first decimal place and Vendor 2 to the second decimal place. A casual check of dozens of randomly-selected short trips in Manhattan from Vendor 1 against their GoogleMaps shortest distances makes us believe that this meter vendor may have rounded down the actual trip distance. For example, a trip of 1.05 miles could be recorded as 1 mile, and this measurement error can make taxi trips appear to be shorter than matched Uber trips by a sizable margin. Moreover, the bias is likely to be larger for short routes than for long routes in percentage terms. Therefore, we drop matches involving taxi trips reported by Vendor 1 and add them back to the sample only as a robustness check.

Table 1: Summary Statistics

Variable	**Mean**	**Std.Dev.**	**10th**	**Median**	**90th**
Taxi trip distance(miles)	8.69	5.32	1.12	9.67	16.55
Uber trip distance(miles)	8.43	5.34	1.14	9.15	16.66
Taxi distance/ Uber distance	1.04	0.16	0.86	1.01	1.26
Taxi trip duration(minutes)	28.65	16.70	7.98	27	50.53
Uber trip duration(minutes)	30.03	17.30	8.72	28.32	52.77
Taxi duration/ Uber duration	0.97	0.22	0.71	0.95	1.25
Airport	0.71	0.45	0	1	1
LaGuardia	0.60	0.49	0	1	1
JFK	0.10	0.30	0	0	1
Newark	0.02	0.12	0	0	0
Non-local passenger	0.51	0.50	0	1	1
Surge multiplier	0.11	0.27	0	0	0.50
Uber driver total trips	2489.19	2008.73	358	2018	5313
Uber driver lifetime rating	4.75	0.09	4.64	4.76	4.85
Uber rider total trips	115.34	170.12	5	55	293
N	95,357				

Notes. Sample is from taxi and Uber trip records of NYC, Jan to June 2016. "Non-local passenger" takes the value 1 if the billing zip code of a given passenger is outside of NYC, or in the case of missing billing zip codes, the passenger's city for most Uber trips is not NYC.

Using the 2016 taxi and Uber data, the matching process generates a sample of 95,357 pairs of matched trips that contains trips from 23,974 Uber drivers. The sample is summarized in Table 1. An average route is about 8.4-8.7 miles long, 28-30 minutes in trip duration. There is a considerable amount of dispersion in both trip distance and trip duration, because our sample is over represented by airport trips and short, within Manhattan trips, compared to the population of taxi trips and Uber trips, as a result of the matching criteria. As the taxi-Uber trip distance ratios and duration ratios suggest, despite starting and ending at the same locations, taxi trips are on average slightly longer in distance and slightly shorter in duration than their matched Uber trips. "Non-local passenger" takes the value 1 if the billing zip code of a given Uber passenger is outside of NYC, or in the case of missing billing zip codes, the passenger's city of most Uber trips is not NYC. The Uber surge multiplier is on average 1.11, and we subtract 1 from surge multipliers for easier interpretation of estimation results later. Therefore, a surge multiplier 0 in our sample means a "1x", or just base fare. An average Uber driver in the sample has driven 2,489 trips before the current trip, and this measure has a large variation. Like other digital platforms, Uber driver lifetime rating is highly concentrated with a mean of 4.75. Lastly, Uber riders in the sample on average have taken 115 Uber trips.

The cumulative distributions of Taxi-Uber distance ratios and duration ratios are separately plotted in Figure 4, across three route types: routes between JFK and Manhattan where taxi fares are fixed, non-airport routes, and all other airport routes where both taxi and Uber fares are metered, which include all LaGuardia trips, all Newark trips, and trips between JFK and NYC outer boroughs. The graph reveals several interesting facts. Figure 4a shows that taxi and Uber trips are quite similar in trip distance for non-airport trips and JFK flat-fare trips, as illustrated by the high concentration of the distance ratios around 1. However, taxi trips are significantly longer in distance than matched Uber trips for all other airport trips, as indicated by the location and shape of the cumulative distribution function. On the other hand, Figure 4b shows that taxis overall arrive faster than Uber, which is also consistent with the implications of our theory. In particular, these patterns appear to reflect driver moral hazard, where taxi drivers seem to be less efficient in routes that are anecdotally more lucrative and driving at a greater average speed than comparable Uber drivers. In the next section, we turn to formal tests of driver moral hazard.

Figure 4: Distance and Duration Ratios of Matched Taxi and Uber Trips

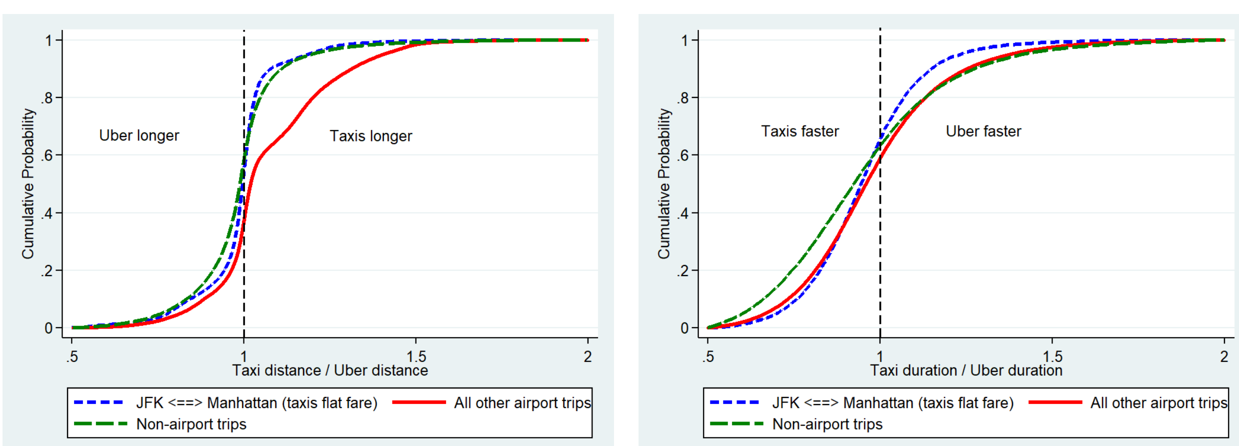

(a) Distance ratios of matched taxi and Uber trips

(b) Duration ratios of matched taxi and Uber trips

5 Empirical Analysis

5.1 Empirical Model

In this section, we use the theory predictions to motivate hypothesis development. Equation (1) requires that for a given matched taxi-Uber pair,

$$\frac{d^c - d^u}{\underline{d}} = \Delta d^c - \Delta d^u + x^{c*} - x^{u*} + \epsilon_d^c - \epsilon_d^u. \tag{6}$$

That is, the normalized difference in Uber and taxi routing is a function of driver skills, detours, and driver-trip random shocks, where the difference in detours is a function of various route characteristics that affect driver incentives.

According to comparative statics in Section 3, both taxi and Uber driver should have little incentive to detour on short trips in thick markets (e.g. trips start and end in Manhattan), because of low marginal payoff of detour (due to short length) and high opportunity cost of detour (finding another ride at drop-off is easy). To add empirical support to this assumption, we investigate how taxi and Uber driver routing in non-airport trips of our sample compares with a measure of long-run average optimal routing given by GURAFU, Uber's internal routing engine. Shown in Figure 5, the GURAFU-weighted trip distance is concentrated to 1 to a large degree, for both taxis and Uber. We consider this to be empirical support for our above-mentioned assumption and take non-airport trips as our control groups in the estimation.

Figure 5: Taxi and Uber driver routing compared with Uber's internal routing engine

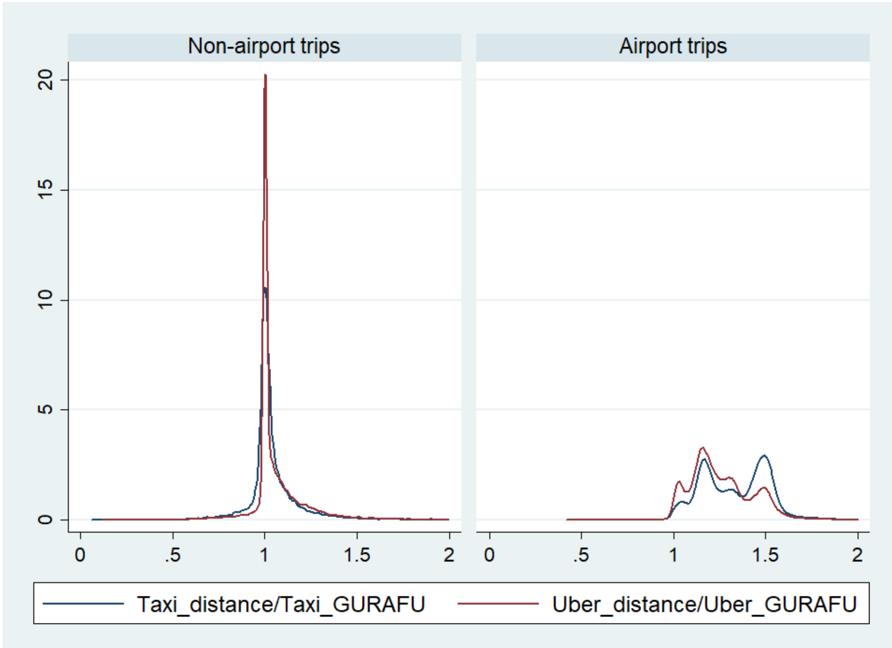

Let r denote a given matched taxi-Uber pair. Based on the theory predictions and the data structure, we specify and estimate the following empirical model on the sample that consists of non-airport trips and airpor trips that are both metered by taxis and Uber (that is, we exclude JFK flat-fare routes):

$$\begin{aligned}
\frac{d_r^c}{d_r^u} &= \alpha_0 + \alpha_1 Airport_r + \alpha_2 Non_local_r + \alpha_3 Airport_r \times Non_local_r \\
&+ \alpha_4 Uber_surge_multiplier_r + \alpha_5 Airport_r \times Uber_surge_multiplier_r \\
&+ \alpha_6 log(Uber_driver_total_trips_r) + \alpha_7 Uber_driver_rating_r + \alpha_8 log(Uber_rider_total_trips_r) \\
&+ \phi_{hw} + \eta_i + \epsilon_{dr}^c - \epsilon_{dr}^u,
\end{aligned} \tag{7}$$

First of all, we include in the regression a set of route characteristics that affect the incentives and costs of detour, which include airport dummy, non-local dummy, airport times non-local, Uber surge multiplier, and airport times Uber surge multiplier. Although detour incentives increase for both taxi and Uber drivers in situations such as longer routes and routes with non-local passengers, Uber's incentive design adds extra penalties to detours.

The effect of these incentives can be reflected in the estimates of these route characteristics. Therefore, we hypothesize that $\alpha_1 > 0$, $\alpha_2 > 0$, $\alpha_3 > 0$, $\alpha_4 < 0$, and $\alpha_5 < 0$.

Second, a set of Uber driver and Uber rider characteristics are controlled for to further explain the variation in Uber-taxi routing difference. Uber driver experience is measured by the driver's total trips driven prior to the current trip, and it is expected to positively correlate with Uber driver routing efficiency if there is a learning-by-doing effect. Uber driver routing efficiency is expected to positively correlated with Uber driver rating as routing efficiency is an important metric in overall driver quality. We also account for Uber rider experience on the platform, measured by the total number of trips completed, as more experienced riders may make the trip more efficient by better communication with the driver, choosing a more efficient pickup/dropoff location, etc. Let ϕ_{hw} represent fixed effects at the hour-of-week level. In some specifications, we include Uber driver fixed effects η_i to control for Uber driver-invariant unobservables.

We can perform a similar decomposition as the one in Equation 6 to Uber-taxi travel time difference:

$$\frac{t^c - t^u}{\underline{t}} = \Delta t^c - \Delta t^u + \lambda(x^{c*} - x^{u*}) + y^{c*} - y^{u*} + \epsilon_t^c - \epsilon_t^u. \tag{8}$$

As discussed in the first-order conditions of the model, taxi drivers overall are incentivized to travel at a greater speed and arrive faster than Uber drivers. To the extent that the total travel time is affected by both the amount of detour and additional speeding, the Uber-taxi difference in travel time should be driven by factors that affect the routing difference $(x^{c*} - x^{u*})$, as well as factors that affect the speeding difference $(y^{c*} - y^{u*})$. Therefore, we estimate a similar empirical model as Equation 7,

$$\frac{t_r^c}{t_r^u} = \beta_0 + \beta_1 Airport_r + \beta_2 Non\_local_r + \beta_3 Airport_r \times Non\_local_r$$
$$+ \beta_4 Uber\_surge\_multiplier_r + \beta_5 Airport_r \times Uber\_surge\_multiplier_r$$
$$+ \beta_6 log(Uber\_driver\_total\_trips_r) + \beta_7 Uber\_driver\_rating_r + \beta_8 log(Uber\_rider\_total\_trips_r)$$
$$+ \phi_{hw} + \eta_i + \epsilon_{tr}^c - \epsilon_{tr}^u. \tag{9}$$

Here we face somewhat competing hypotheses. $\beta_1 > 0$ if for airport trips, taxi additional travel time because of detour exceeds the travel time saved by speeding, rendering the total travel time to be greater than that of Uber; $\beta_1 < 0$ if the opposite is true. A similar logic applies to β_2 and β_3. However, it is unambiguous to expect that $\beta_4 < 0$ and $\beta_5 < 0$, because when surge multiplier is greater, it is in Uber drivers' best interest to detour and not to speed.

5.2 Identification

Without loss of generality, we assume normality of the mean-zero driver-route random shocks ϵ_{dr}^u, ϵ_{dr}^c, ϵ_{tr}^u, and ϵ_{tr}^c. For Ordinary Least Squares estimation to produce unbiased coefficient estimates, route characteristics in the empirical model need to be uncorrelated with unobserved driver-route shock in the error term. This is supported by the institutional features of taxis and Uber.

On one hand, for taxi drivers, the matching of passengers of certain destinations is close to randomly assigned, because passengers determine the destinations and refusals of passengers are heavily penalized by the TLC refusal law.[7] However, taxi drivers can indeed form expectations of passenger destinations and route profitability by developing their own search strategies, leading to a correlation between route characteristics and driver types. In

[7] Per the TLC refusal law, "It is against the law to refuse a person based on race, disability, or a destination in New York City. A taxicab driver is required to drive a passenger to any destination in the five boroughs." Riders are encouraged to make a refusal complaint by calling 3-1-1. According to Haggag et al. (2017), "In 2009 the refusal punishment was \$200-\$350 for a first offense, \$350-\$500 and a possible 30-day license suspension for a second, and a mandatory license revocation for a third offense. The TLC received about 2,000 formal complaints per year in 2009 and 2010"

this case, taxi driver fixed effects are a good way to tease out the bias. however, it is infeasible in the main analysis with 2016 taxi data due to the lack of taxi driver IDs. Nonetheless, we demonstrate in Section 6.1 that driver selection appears to be insignificant when the same estimation is run on 2013 data where taxi driver fixed effects are controlled for.

On the other hand, several features of the Uber platform limit the scope of endogeneity of route characteristics:

1. To Uber drivers, rider assignment by the platform is virtually random by construction. Uber's matching of drivers and riders are mainly based on spatial proximity and dispatching efficiency, and it gives little weights to driver and rider characteristics in the matching. Therefore, route characteristics can be viewed as exogenous.

2. Uber drivers have the option to cancel trip requests, but cancellation of rides is costly. If drivers could cancel on riders without any costs, they would do it in order to select more profitable rides, thus creating an endogeneity problem. However, trip cancellation is constrained at various places on Uber. Once assigned a rider, the driver cannot see the rider's destination on the application until having picked up that rider, which makes it difficult for drivers to "cherry pick" passengers before accepting a trip request. Moreover, frequent and suspicious ride cancellation is penalized on Uber, often in the form of warning, "time out", or even deactivation. In addition, it is difficult for a driver to form expectations on the next rider' profitability after cancellation, making cancellation of the current ride risky.

3. For individual Uber drivers, surge pricing is difficult to predict and chase. As shown in Lam and Liu (2017), Uber surge multipliers are extremely volatile and hard to predict by highly granular location-time fixed effects. With such volatile and unpredictable surge multipliers, it is not in the driver's best interest to chase the surge, at the cost of forgone earnings from trip requests declined. Uber drivers commonly agree on this view, based on our conversations with Uber drivers in New York City and Boston, Massachusetts. In addition, the Uber app no longer shows the surge hot spots after driver's accepting a ride, which prevents drivers from canceling the current no-surge or low-surge ride in order to get a high-surge ride.

Taken together, these aforementioned institutional details suggest that the correlation between route characteristics and unobserved driver-route shocks is at best limited. To further reduce the potential bias, we control for Uber driver fixed effects in some specifications.

5.3 Results

5.3.1 Baseline: Non-airport Routes + Metered Airport Routes

The main results of our analysis are shown in Table 2, which are performed on all routes that are both metered on Uber and taxis (that is, all routes except JFK taxi flat-fare routes). Specifications (1)-(3) show regression results of taxi-Uber distance ratios, and Specifications (4)-(6) show regression results of taxi-Uber duration ratios.

In Specification 1, we find a sizable and statistically strong effect of airport trips on the relative size of taxi-Uber detour ratio. Specifically, the taxi-Uber distance ratio for an airport trip is on average 7.4% larger than for a non-airport trip, other things held constant. The effect of non-local passengers is small and weakly estimated on non-airport trips, yet its effect is positive and strong on airport trips, suggesting that taxi drivers tend to exploit the information advantage more on profitable routes. However, it is important to note that we use the information of the Uber rider to proxy for whether the taxi passenger is a local or not. Therefore, this variable could be subject to measurement errorrrors, which may lead to a downward bias when it is a classical measurement error. We caution that the scope of measurement error of "non-local" should be smaller on airport routes than on non-airport routes, because it is more likely that the taxi rider and the Uber rider are either both locals or both non-locals when they head to the airport from the same specific place at the same time (eg. a hotel). Lastly, we find that Uber drivers tend to detour more relative to taxi drivers on airport trips with a surge price in effect. However, the effect of surge multiplier is not found on non-airport trips. In addition, we find that Uber driver rating is positively correlated with the relative detour, suggesting that Uber drivers with better ratings are more efficient at routing or drivers who take inefficient routes are more likely to get low ratings.

Table 2: Taxi and Uber Driver Routing Difference

	D.V. = Taxi dist./Uber dist.			D.V. = Taxi dur./Uber dur.		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Airport	0.074***	0.074***	0.074***	0.034***	0.038***	0.036***
	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Non_local	-0.002	-0.003	0.001	0.008**	0.008**	0.006
	(0.002)	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)
Airport × Non_local	0.018***	0.019***	0.013***	-0.009**	-0.011***	-0.008*
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)
Surge_multiplier	0.002	0.001	-0.000	0.008*	0.009*	-0.002
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.006)
Surge_multiplier × Airport	-0.026***	-0.023***	-0.022***	-0.021***	-0.013*	-0.006
	(0.005)	(0.005)	(0.005)	(0.008)	(0.007)	(0.009)
Log (Uber_driver_total_trips)	0.000	-0.000	0.005**	0.010***	0.010***	0.012***
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.003)
Uber_driver_rating	0.040***	0.039***	0.000	0.128***	0.126***	0.000
	(0.007)	(0.007)	(.)	(0.009)	(0.009)	(.)
Log (Uber_rider_total_trips)	0.001	0.001	0.001	0.002***	0.003***	0.003***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Hour of week FE	No	Yes	Yes	No	Yes	Yes
Uber driver FE	No	No	Yes	No	No	Yes
N	86,627	86,627	86,627	86,627	86,627	86,627
R^2	0.055	0.060	0.376	0.008	0.014	0.330

Notes. For all specifications, standard errors are cluster-robust at the hour-of-week level.

*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

The estimated effects hardly change when hour-of-week fixed effects in Specification (2). Further controlling for Uber driver fixed effects does not appear to affect the point estimates much, as shown in Specification (3). Interestingly, the effect of Uber driver total trips becomes strong, suggesting a routing improvement due to accumulated driving experience within a driver.

In Specifications (4)-(6), the same set of regression analyses are performed to explain the variations in taxi-Uber travel time difference. Across specifications, we observe a positive and strong effect of airport trips. Using Specification (6) as an example, the taxi-Uber travel time ratio of an airport trip is on average 3.6% greater than that of a non-airport trip, other things held constant. Combining this effect with the 7.4% effect of airport routes on distance ratio suggests that (1) taxi drivers are more additionally faster than Uber drivers on airport routes than on non-airport routes, and (2) this additional speeding may not fully compensate the time spent on detour. The effect of "non-local" is positive on non-airport routes and the effect disappears on airport trips (0.006 - 0.008 = 0.002 using Specification (6)). Across specifications, the effects of "non-local" are estimated with noise and small in size. Similarly, the effects of surge multiplier are also weakly identified, although overall they exhibit expected signs. Contrary to the routing distance, routing time is much more affected by driver experience and driver rating. In particular, driver experience affects travel time efficiency both across drivers (the extensive margin) and within drivers (the intensive margin). Thus it appears that while the technology tools and incentive design that the Uber platform bring to bear are important, driver expertise and learning also continue to have important roles.

5.3.2 Non-airport Routes + JFK Taxi Flat-fare Routes

It is important to note that the analysis so far identifies moral hazard in relative terms — the effect of 7.4% describes how taxi driver routing in metered airport routes, referenced by Uber driver routing in these routes, compares with that in non-airport routes. Therefore, taxi drivers may in fact detour more in absolute terms if the comparable Uber drivers also detour to some extent. In this section, we leverage the JFK taxi flat-fare routes to identify the extent to which Uber drivers detour — these routes are an ideal scenario where taxi drivers do not have any incentive to detour when the trip fare is flat, because detour only incurs cost while not gaining any benefit. In particular, we estimate the same regressions as Equation 7 and Equation 9 on the sample that consists of non-airport routes and routes between Manhattan and JFK.

The regression results are shown in Table 3. The effect of airport trips(hereby JFK trips)

Table 3: Taxi and Uber Driver Routing Difference: JFK Taxi Fixed-fare Routes

	D.V. = Taxi dist./Uber dist.			D.V. = Taxi dur./Uber dur.		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Airport	0.003	0.002	-0.002	0.005	0.012***	0.015**
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.007)
Non_local	-0.002	-0.002	-0.002	0.008**	0.008**	0.007
	(0.002)	(0.002)	(0.003)	(0.004)	(0.004)	(0.005)
Airport × Non_local	0.001	0.002	0.006	-0.006	-0.009*	-0.013
	(0.003)	(0.004)	(0.005)	(0.005)	(0.005)	(0.008)
Surge_multiplier	0.002	0.000	-0.001	0.008*	0.008*	-0.009
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.008)
Surge_multiplier × Airport	-0.007	-0.008	-0.003	-0.024***	-0.025***	-0.005
	(0.008)	(0.008)	(0.013)	(0.009)	(0.009)	(0.015)
Log (Uber_driver_total_trips)	0.002***	0.002***	0.000	0.008***	0.008***	0.015***
	(0.001)	(0.001)	(0.003)	(0.001)	(0.001)	(0.005)
Uber_driver_rating	0.029***	0.028***	0.000	0.137***	0.136***	0.000
	(0.008)	(0.008)	(.)	(0.013)	(0.013)	(.)
Log (Uber_rider_total_trips)	0.001	0.001	0.000	0.001	0.001	0.002
	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Hour of week FE	No	Yes	Yes	No	Yes	Yes
Uber driver FE	No	No	Yes	No	No	Yes
N	36,587	36,587	36,587	36,587	36,587	36,587
R^2	0.001	0.007	0.520	0.004	0.011	0.522

Notes. For all specifications, standard errors are cluster-robust at the hour-of-week level.
*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

on the taxi-Uber distance ratio is not statistically different from zero, which suggests that Uber drivers do not detour on routes between JFK and Manhattan. Similarly, Uber drivers do not route less efficiently when the passenger is non-local or when the surge multiplier is high, for both airport routes and non-airport routes. However, the effect of airport on the duration ratio is strongly estimated in Specifications (5) and (6), although the effect is small in size. This is consistent with the patterns in Figure 4b, where taxi-Uber speed difference is even larger in non-airport trips than in airport trips. It is likely because taxi drivers have a greater incentive to speed in busy location-hours to avoid the low rate by minute ($0.5 per minute), which is lower than if they driver above 12 mile per hour to earn by mileage. Effects of non-local and surge multiplier are generally small and weakly estimated, with the exception of the interaction of surge multiplier and airport in Specification (4) and (5), meaning that Uber drivers tend to drive at a slower pace on airport routes with high surge prices than they do on non-airport routes. This is consistent with what we find in the baseline results in Section 5.3.1.

Regression results presented in this section complement the baseline results, where we show that the observed driver strategic inefficiency is indeed the true size of moral hazard. Combining these two sets of analyses gives rise to our main empirical findings: 1. taxi drivers detour on airport trips that are metered; 2. taxi drivers overall drive faster than Uber drivers; 3. Uber drivers generally do not detour or strategically drive slow, except on airport trips with high surge prices.

6 Competing Explanations and Robustness

In this section, we examine three alternative explanations, and conclude that none of these is likely to have caused the observed patterns in taxi-Uber routing difference.

6.1 Driver Selection

As discussed in the identification section, the unobserved taxi and Uber driver types in the error term may be correlated with route characteristics. For example, strategic driver types may more likely take airport routes and routes more traveled by non-local riders. If this is the case, then the observed effects are in fact largely driven by driver selection into certain routes, rather than drivers being responsive to incentives. It is revealing that we see no significant changes in the coefficient estimates when Uber driver fixed effects are controlled for in the main analysis. However, we need to explore whether the effects remain when taxi driver fixed effects are accounted for.

We repeat the same regression analysis on the 2013 data. Given the small market share of Uber in 2013, matches that follow the same matching procedure as for 2016 data would lead to a sample size too small for identification, especially when the aim is to purge out the within-taxi-driver variations. Therefore, we relax the matching criteria to only Step 1 and Step 4, i.e. we do not restrict the matched trips to be on the same street, or following the same traffic direction. Also, we relax the time window to be 30 minutes. In addition, a small share of matches are dropped because new Uber drivers did not have ratings at the time of the trip. The final sample consists of 16,989 matches on metered routes with 10,085 taxi drivers, 4,410 of which have more than one trip in the sample.

Shown in Table 4, estimates are in general of the same sign and similar size with that of 2016 analysis. Particularly, in Specification (3), the effect of airport on distance ratio remains strong and large when taxi driver fixed effects are controlled, and the effect of non-local on airport routes is borderline significant. For travel time difference, the estimates lose statistical significance when taxi driver fixed effects are added in Specification (6). Overall, the findings based on 2013 data suggest that behavioral responses to incentives and opportunities are present, and driver selection is not fully responsible for the observed routing difference.

It is also possible that drivers of different types select differently into being taxi and Uber drivers. If this is the case, then the observed moral hazard can be an artifact of the driver type distributions of taxis and Uber. Not being able to directly observe driver types, we cannot definitively rule this possibility out. However, we shed light on the extent of driver behavioral change by following 1,549 former taxi drivers who switched to Uber — given that types are persistent, former taxi drivers who used to detour might be expected to continue their detour behavior on Uber in similar situations. Table 8 shows the regression results of distance ratios and duration ratios on the sample of non-airport routes and metered airport routes, for these 1,549 former taxi drivers and their matched Uber drivers in 2013. The estimates confirm the moral hazard behavior of these former taxi drivers in 2013. Table 9 shows the regression results of distance ratios and duration ratios on the sample of non-airport routes and JFK flat-fare routes, for these drivers and their matched taxi drivers in 2016. Results here show little evidence that these drivers continue to detour in profitable routes, when compared with the behavior of Uber drivers. This provides strong evidence that drivers adapt to changes in technology and market design via behavioral change, which is consistent with the driver moral hazard finding of this paper.

6.2 A GPS Story?

It is also possible that the Uber-taxi routing difference is driven by the difference in technology adoption — Uber drivers generally adopt GPS navigation while NYC taxi drivers mostly rely on their own judgment. For example, the effect of airport can also be rationalized if the

Table 4: Taxi and Uber Driver Routing Difference, 2013

	D.V. = Taxi dist./Uber dist.			D.V. = Taxi dur./Uber dur.		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Airport	0.105***	0.112***	0.101***	0.017**	0.013	0.003
	(0.007)	(0.007)	(0.011)	(0.008)	(0.008)	(0.017)
Non_local	-0.007**	-0.010***	-0.009	0.001	0.002	-0.006
	(0.003)	(0.003)	(0.006)	(0.005)	(0.005)	(0.010)
Airport × Non_local	0.019***	0.022***	0.023*	-0.007	-0.008	0.002
	(0.007)	(0.007)	(0.013)	(0.010)	(0.010)	(0.019)
Surge_multiplier	-0.001	-0.002	-0.006	0.006	0.009	0.001
	(0.006)	(0.004)	(0.007)	(0.006)	(0.006)	(0.012)
Surge_multiplier × Airport	-0.027**	-0.021**	-0.035	-0.028*	-0.027	-0.016
	(0.011)	(0.010)	(0.024)	(0.016)	(0.017)	(0.038)
Log (Uber_driver_total_trips)	-0.002**	-0.003**	-0.003	0.001	0.000	-0.002
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)
Uber_driver_rating	0.056***	0.052**	0.048	0.100***	0.103***	0.077
	(0.020)	(0.021)	(0.039)	(0.031)	(0.031)	(0.056)
Log (Uber_rider_total_trips)	0.000	0.000	0.001	0.002	0.001	0.004
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)
Hour of week FE	No	Yes	Yes	No	Yes	Yes
Taxi driver FE	No	No	Yes	No	No	Yes
N	16,989	16,989	16,989	16,989	16,989	16,989
R^2	0.082	0.102	0.644	0.001	0.015	0.620

Notes. For all specifications, standard errors are cluster-robust at the hour-of-week level.
*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

benefits of GPS over taxi driver knowledge is greater for these routes than on non-airport routes.

Regression results on the sample of JFK flat-fare routes directly speak against this possibility — JFK routes are among the longest routes of NYC, yet no significant difference in Uber and taxi routing is observed on these routes. However, it is likely that taxi drivers are more familiar with JFK routes, compared to other long routes that are infrequently traveled, and as a result, JFK routes are not a conclusive counterargument. Noting this, we split LaGuardia pick-ups into trips to Manhattan and trips to outer boroughs and show in Figure 6 that LGA-to-Manhattan routes, the routes that should be more familiar to taxi drivers, are exactly where taxi drivers route *longer* in distance. Given the large share of LGA trips in our sample, the observed taxi-Uber routing difference cannot reasonably be interpreted as solely a GPS story.

Figure 6: LaGuardia Trips to Manhattan vs. to Outer Boroughs

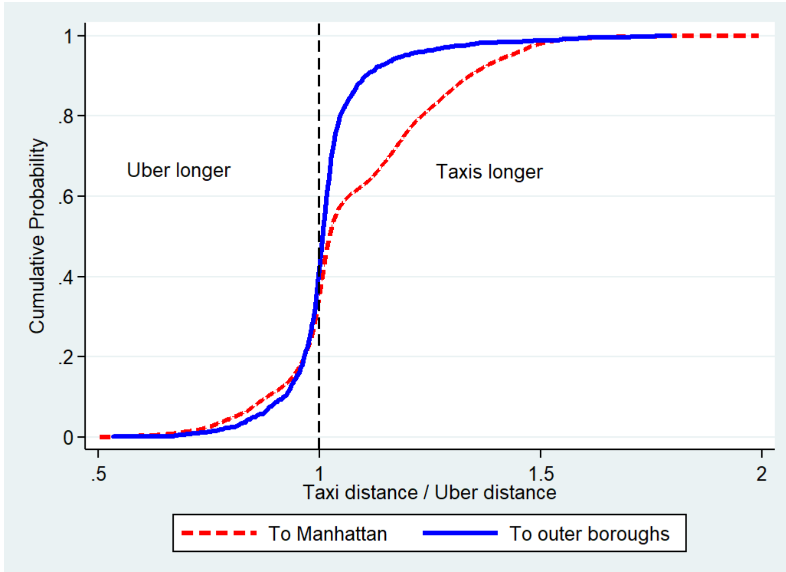

6.3 Taxi Drivers Take Longer But Time-Efficient Routes?

It is also possible that taxi drivers take routes that are longer in trip distance and shorter in trip duration, compared to Uber drivers, which can also lead to the observed effects. Among alternative routes, taxi drivers indeed have incentives to choose a route that is both long in distance and short in duration. But in order for this to be the main explanation of our findings, either Uber GPS systematically prioritizes route distance over travel time, or taxi drivers on average possess superior routing information than Uber GPS at detecting long but time efficient routes. Our conversation with Uber's routing department rules out the first possibility – in fact, Uber GPS weighs travel time more than trip distance. Therefore, we need to investigate whether taxi drivers possess superior information, and a necessary condition for that is that taxi drivers finish trips faster when they take a longer route compared to matched Uber trips.

Figure 7: Taxi-Uber Duration Ratios and Speed Ratios

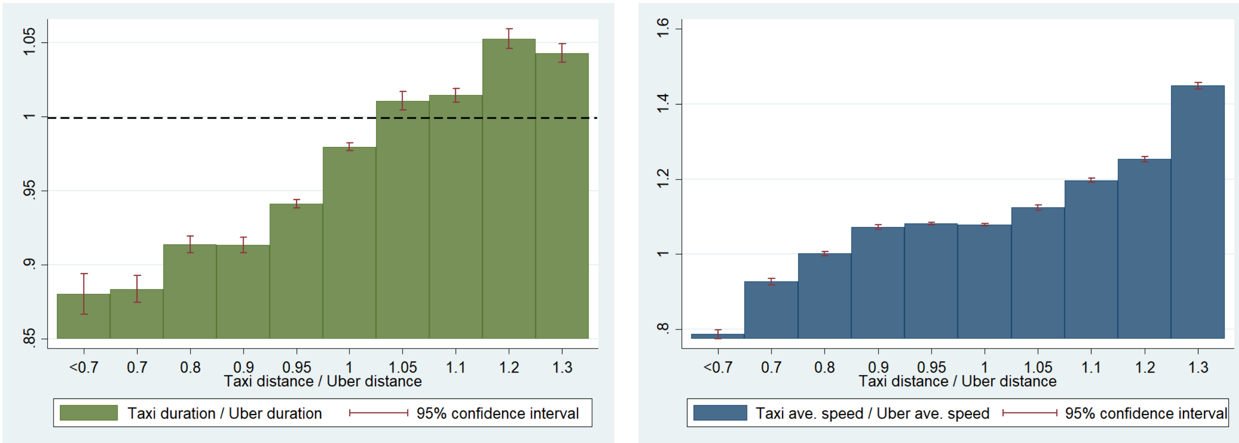

(a) Taxi-Uber duration ratios across taxi-Uber distance ratios

(b) Taxi-Uber speed ratios across taxi-Uber distance ratios

As shown in Figure 7 (a), when detouring (ie. distance ratio $>$ 1), taxi drivers on average finish trips *later* than comparable Uber drivers. Furthermore, the relative travel time difference increases when taxi detour is longer. A look into the relative average travel speed in Figure 7 (b) reveals that taxi drivers tend to speed more when detouring, but the speeding cannot fully make up the time lost in detours. Therefore, we conclude that taxi drivers do not consistently out-perform Uber drivers at detecting long but time-efficient routes.

6.4 Robustness

We perform two sets of robustness checks to strengthen the identification of driver moral hazard. In the main analysis of driver detour, we constrained the matched taxi and Uber trips to be 15 minutes apart. In the first set of robustness checks, we perform the same analysis using alternative time windows, namely 5 minutes, 10 minutes, 20 minutes, and 30 minutes. Shown in Table 6, the estimated effects are stable and consistent across time window lengths, particularly for the regressions with distance ratio as the dependent variable. To the extent that trips within a smaller time range can be more subject to the same traffic and thus better approximate the experimental ideal, the fact we find significant effects and they are of similar size even with a time window as short as 5 minutes greatly enhances our identification.

Recall that our main sample only contains taxi trips reported by Vendor 2, because Vendor 1's meter system appears to round down trip distance to the nearest first decimal place. Taxi trips may appear to be shorter because of the rounding. Consider a pair of matched Uber and taxi trips, where the Uber trip is 1 mile and the taxi trip is 0.95 miles but reported taxi trip length is rounded to 0.9 miles. Then the distance ratio would be 0.9 instead of 0.95, with a downward bias of -0.05. For a 9.95-mile taxi trip rounded to 9.9 miles with a matched 10-mile Uber trip, the downward bias is only -0.005. Thus, the same amount of rounding error leads to proportionately greater downward bias on shorter routes. The implication is that for the moral hazard behavior to be consistent across drivers of both vendors, we expect the robustness check on Vendor 1 sample to yield an upward bias in the coefficient estimate of airport trips, instead of the opposite. In Table 7, we separately estimate the main regressions using Vendor 1 only and both Vendor 1 and Vendor 2, and compare these with the main regression results using only Vendor 2. We indeed find the upward bias, where the estimated effect of airport on distance ratio is 9.3% on Vendor 1 sample, 1.9% more than the effect on Vendor 2 sample. We find this upward bias, instead of a downward one, consistent with our main findings.

7 Discussion

7.1 Mechanisms

In this section, we discuss two mechanisms that account for the observed routing difference between taxi and Uber drivers.

The first mechanism is the set of technology-enabled incentive devices implemented by the Uber platform but not by taxis. These incentive devices include tech-aided monitoring and verification, tech-enabled rider rating of drivers, and tech-aided conflict resolution. Each of these make the cost function for moral hazard via detour steeper for Uber than for taxis. One necessary condition for a working rating system to penalize strategic behavior is the negative correlation between passengers' ratings to the drivers and driver routing inefficiency. To test this correlation, we use the subsample of only JFK taxi flat-fare routes as taxi driver routing in these routes is a valid benchmark. Shown in Table 5 (1) - (5), both the normalized Uber distance (by taxi distance) and the normalized Uber duration (by taxi duration) negatively predict the ratings given by passengers, conditional on the fact that a rating was given. The correlations is more salient in the subset where Uber distance is greater than taxi distance (Specification (6)). However, in the cases where Uber distance is less than taxi distance (Specification (7)), Uber trips shorter in distance than the matched taxi trips are less likely to gain high ratings. One likely reason for this is that passengers dislike off-GPS routing and thus give low ratings even when Uber drivers have found a shorter route. It seems plausible that passengers cannot easily assess whether a driver's deviations from the prescribed GPS route are due to superior information used to shorten the route or an effort to extract a higher fare with a longer route.

These incentive devices all appear to enhance market transparency and our results suggest they are effective at mitigating moral hazard in this setting. Our results also imply that these mechanisms govern most cases as it is seen that Uber drivers do not appear to detour on JFK routes, where the potential gain from detour is large. Strikingly, Uber drivers do tend to detour or driver more slowly on airport trips when surge pricing is high, suggesting that the the monitoring devices are not binding. Instead, as the agency model predicts, Uber drivers are willing to make a trade off when surge pricing is high that can offset possible penalty.

In addition to technology-enabled monitoring tools, pricing is another important mechanism that predicts driver routing behavior. This is most clearly reflected in taxi driver routing efficiency on JFK routes with flat fares. When taxi fares are metered as a two-part tariff, taxi drivers tend to detour on longer routes because the variable part of the fare can justify the detour; on the other hand, taxi drivers tend to detour less in short routes, especially in thick markets like Manhattan, because it is in their best interest to take as many trips as possible to exploit the proportionately larger fixed component of the fare.

Table 5: Uber Driver Ratings by Passengers Are Correlated with Routing Efficiency

	All JFK taxi flat fare routes					Uber dist. ≥ Taxi dist.	Uber dist. < Taxi dist.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Uber_dist/Taxi_dist	-0.222***		-0.152**	-0.150**	-0.132	-0.283**	0.518**
	(0.061)		(0.062)	(0.061)	(0.082)	(0.115)	(0.218)
Uber_dur/Taxi_dur		-0.276***	-0.253***	-0.204***	-0.206***	-0.269***	-0.123
		(0.045)	(0.046)	(0.046)	(0.053)	(0.077)	(0.082)
Uber_driver_rating				0.862***	0.880***	0.904***	0.765***
				(0.096)	(0.104)	(0.149)	(0.154)
Log (Uber_driver_total_trips)				0.013*	0.015*	0.021	0.011
				(0.008)	(0.009)	(0.013)	(0.012)
Surge_multiplier				-0.111**	-0.100*	-0.121	-0.034
				(0.045)	(0.053)	(0.083)	(0.079)
Non_local				0.019	0.021	0.039	0.007
				(0.017)	(0.023)	(0.031)	(0.030)
Log (Uber_rider_total_trips)				0.028***	0.027***	0.035***	0.018**
				(0.005)	(0.006)	(0.009)	(0.008)
Hour of week FE	No	Yes	Yes	No	Yes	Yes	Yes
N	5,885	5,885	5,885	5,885	5,885	3,139	2,746
R^2	0.002	0.006	0.007	0.026	0.049	0.080	0.078

Notes. For Specifications (5), (6), and (7), standard errors are cluster-robust at the hour-of-week level. *** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

Figure 8: Driver Marginal Minute Earning Across Travel Speeds

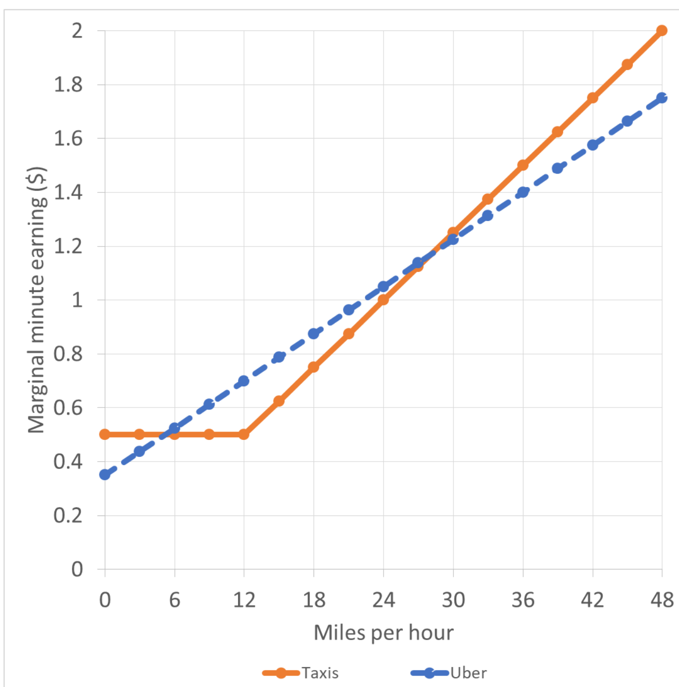

Furthermore, perhaps the pricing scheme is of first-order importance in explaining taxi driver travel speed. We document in the main analysis that taxi drivers, although detouring in some cases, in general drive at a faster speed than Uber drivers. This is expected because taxi drivers are not paid by their driving time in normal traffic. Next, we compare the pricing schemes of taxis and Uber and show that speeding is rewarded more for taxi drivers than for Uber drivers. In Figure 8, we compute driver marginal minute earnings at various traffic speeds, separately for taxis and Uber, following their pricing formulas. All the computations are net of the fixed component of the fare (i.e. \$2.5 for taxis and \$2.55 for Uber). An interesting divergence appears. For example, taxi driver per-minute earning increases by 4 times from 12 MPH to 48 MPH (from \$0.5/min to \$2/min), while an Uber driver per-minute earning only increases by 2.5 times (from \$0.7/min to \$1.75/min). This difference stems from the weight given to trip distance in the pricing formula, where taxi distance is marginally more rewarding than that of Uber. One artifact of the taxi pricing schedule is that, when traffic is flowing at about 12 miles per hour, a NYC taxi driver would actually earn a slightly higher fare by alternating between being stopped and racing ahead at 24 miles per hour.

7.2 Mind vs. Machine

"GPS routes are slower in Manhattan, so I don't use it."

— Gurpreet Singh, NYC taxi driver, interviewed on July 9th, 2018.

The data show that taxi drivers are efficient at routing than Uber drivers in short, non-airport routes. For routes within Manhattan Core, where Manhattan Core is defined as roughly the part of Manhattan below the north edge of Central Park, taxi trip distance is on average 98.5% of that of Uber, and the difference is statistically significant at 1% level. This is evidence that human navigation still can perform at least as well as the technology in dense markets.

One way that human navigation can outperform GPS is by having more up-to-date information on the road networks and conditions, for example, temporary road closures, upcoming sporting events, or undocumented short-cuts. Another possibility is that experienced taxi drivers can suggest a better drop-off point than the exact address given by the passenger, based on the driver's extensive experience. For example, the driver might suggest dropping off the passenger on the opposite side of the street in order to avoid unnecessary travel. This is confirmed by an interview with Loai Yousef, an NYC Uber and Lyft driver, who stated that "Sometimes the Uber GPS map has mistakes. Sometimes it makes the driver do a U-turn to arrive at the exact address even though it would be easy for rider to just cross the street. Taxis drivers often drop passengers off a short distance from exact address."

NYC taxi driver expertise in routing should not be surprising, because as residual claimants, they are strongly motivated to learn the routes, optimize their routing, and take initiative when they can. In contrast, Uber drivers might not have as large of a motivation to use their discretion, even in cases when they do possess better information than the GPS. The reason is that off-GPS routing might come across as suspicious behavior to the riders, which can result in bad ratings and complaints. Loai Yousef shared a similar insight from a practitioner's perspective: "..., Uber passengers tend to want driver to go to the exact address even if it's wasteful." Therefore, the use of GPS, coupled with the monitoring and rating systems, can limit the incentives for human knowledge accumulation, as well as initiative and discretion.

8 Conclusion

In this paper, we study how a technology platform affects agent monitoring and reward systems, and consequently moral hazard and service quality. We provide causal evidence from the taxi and Uber setting in the form of driver choices that affect the length and duration of trips from the identical start and end points. Analyzing trip-level data from NYC, we find that taxi drivers tend to detour more relative to Uber drivers on airport routes and especially routes taken by non-local riders, while Uber drivers tend to detour more relative to taxi drivers on routes with greater surge pricing. In other words, as predicted by our theoretical model, the Uber technology platform and pricing scheme reduce driver moral hazard behavior where taxi moral hazard return is high, but at the same time create other margins of driver moral hazard. What's more, the the incentives for creating and using driver routing expertise may be reduced by the Uber platform. These findings are consistent with the agency theory, while driver selection or difference between taxis and Uber in navigation technologies are found not likely to have caused the observed patterns in driver routing.

We highlight the importance of incentive devices as well as pricing schemes as the underlying mechanisms. As such, our findings extend important implications to industry participants. For TLC, our results provide support for the development and implementation of smart phone applications that handle taxi dispatching and matching with passengers, digital payment, and passenger monitoring. Also, it is important for TLC to re-evaluate the current pricing scheme that rewards taxi cab speeding as well as impacts of alternative pricing structures. For digital platforms such as Uber, our findings suggest an opportunity for machine learning based techniques to detect driver opportunistic behavior, which may further enhance market transparency and trust building.

This paper contributes to the debate over the digital disruption, particularly the part involving digital platforms that represent new labor contractual relationships. Previous studies have demonstrated efficiency gains of ride-hailing platforms due to the ability to better align demand and supply, compared to regulated taxis. Building on these important margins, we identify the efficiency gain from the agency perspective. The efficiency gain is large, given the sizable estimated effects. However, a full welfare account calls for future empirical work such as quantifying the impact of speeding on public safety and traffic congestion.

References

Angrist, J. D., S. Caldwell, and J. V. Hall (2017). Uber vs. taxi: A drivers eye view. Technical report, National Bureau of Economic Research.

Aral, S., E. Brynjolfsson, and L. Wu (2012). Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science 58*(5), 913–931.

Bakos, J. Y. (1997). Reducing buyer search costs: Implications for electronic marketplaces. *Management science 43*(12), 1676–1692.

Balafoutas, L., A. Beck, R. Kerschbamer, and M. Sutter (2013). What drives taxi drivers? a field experiment on fraud in a market for credence goods. *Review of Economic Studies 80*(3), 876–891.

Balafoutas, L., R. Kerschbamer, and M. Sutter (2017). Second-degree moral hazard in a real-world credence goods market. *The Economic Journal 127*(599), 1–18.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics 117*(1), 339–376.

Brown, J. R. and A. Goolsbee (2002). Does the internet make markets more competitive? evidence from the life insurance industry. *Journal of political economy 110*(3), 481–507.

Brynjolfsson, E., Y. Hu, and M. D. Smith (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science 49*(11), 1580–1596.

Brynjolfsson, E. and M. D. Smith (2000). Frictionless commerce? a comparison of internet and conventional retailers. *Management science 46*(4), 563–585.

Castillo, J. C., D. Knoepfle, and G. Weyl (2017). Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 241–242. ACM.

Chen, M. K., J. A. Chevalier, P. E. Rossi, and E. Oehlsen (2017). The value of flexible work: Evidence from uber drivers. Technical report, National Bureau of Economic Research.

Chen, M. K. and M. Sheldon (2016). Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *EC*, pp. 455.

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

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

Duflo, E., R. Hanna, and S. P. Ryan (2012). Incentives work: Getting teachers to come to school. *American Economic Review 102*(4), 1241–78.

Forman, C., A. Ghose, and A. Goldfarb (2009). Competition between local and electronic markets: How the benefit of buying online depends on where you live. *Management science 55*(1), 47–57.

Gans, J. S., A. Goldfarb, and M. Lederman (2017). Exit, tweets and loyalty. Technical report, National Bureau of Economic Research.

Goldfarb, A. and C. Tucker (2017). Digital economics. Technical report, National Bureau of Economic Research.

Guda, H. and U. Subramanian (2018). Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication and worker incentives.

Haggag, K., B. McManus, and G. Paci (2017). Learning by driving: Productivity improvements by new york city taxi drivers. *American Economic Journal: Applied Economics 9*(1), 70–95.

Hall, J., C. Kendrick, and C. Nosko (2015). The effects of ubers surge pricing: A case study. *The University of Chicago Booth School of Business*.

Hall, J. V., J. J. Horton, and D. T. Knoepfle (2017). Labor market equilibration: Evidence from uber. Technical report, Working Paper, 1–42.

Hall, J. V. and A. B. Krueger (2015). An analysis of the labor market for ubers driver-partners in the united states. *ILR Review*, 0019793917717222.

Hubbard, T. N. (2000). The demand for monitoring technologies: the case of trucking. *The Quarterly Journal of Economics 115*(2), 533–560.

Hui, X., M. Saeedi, Z. Shen, and N. Sundaresan (2016). Reputation and regulations: evidence from ebay. *Management Science 62*(12), 3604–3616.

Klein, T. J., C. Lambertz, and K. O. Stahl (2016). Market transparency, adverse selection, and moral hazard. *Journal of Political Economy 124*(6), 1677–1713.

Lam, C. T. and M. Liu (2017). Demand and consumer surplus in the on-demand economy: the case of ride sharing.

Lee, M. K., D. Kusbit, E. Metsky, and L. Dabbish (2015). Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 1603–1612. ACM.

Liang, C., Y. Hong, and B. Gu (2016). Effects of it-enabled monitoring on labor contracting in online platforms: Evidence from a natural experiment.

Liu, T., E. Vergara-Cobos, and Y. Zhou (2017). Pricing schemes and seller fraud: Evidence from new york city taxi rides.

Nagin, D. S., J. B. Rebitzer, S. Sanders, and L. J. Taylor (2002). Monitoring, motivation, and management: The determinants of opportunistic behavior in a field experiment. *American Economic Review 92*(4), 850–873.

Pierce, L., D. C. Snow, and A. McAfee (2015). Cleaning house: The impact of information technology monitoring on employee theft and productivity. *Management Science 61*(10), 2299–2319.

Rajgopal, S. and R. White (2015). Cheating when in the hole: The case of new york city taxis. Technical report, Working Paper.

Reimers, I., B. R. Shiller, et al. (2018). Proprietary data, competition, and consumer effort: An application to telematics in auto insurance. Technical report.

Staats, B. R., H. Dai, D. Hofmann, and K. L. Milkman (2016). Motivating process compliance through individual electronic monitoring: An empirical examination of hand hygiene in healthcare. *Management Science 63*(5), 1563–1585.

Sudhir, K. and D. Talukdar (2015). The peter pan syndrome in emerging markets: The productivity-transparency trade-off in it adoption. *Marketing Science 34*(4), 500–521.

Tabarrok, A. and T. Cowen (2015). The end of asymmetric information. *Cato Unbound*.

Waldfogel, J. (2017). How digitization has created a golden age of music, movies, books, and television. *Journal of Economic Perspectives 31*(3), 195–214.

A Figures

Figure 9: Dividing NYC into Voronoi Cells Centered at Street Intersections

Figure 10: Pickup Locations of Uber and Taxi after Matching Step 1 (taxis in purple; Uber in green)

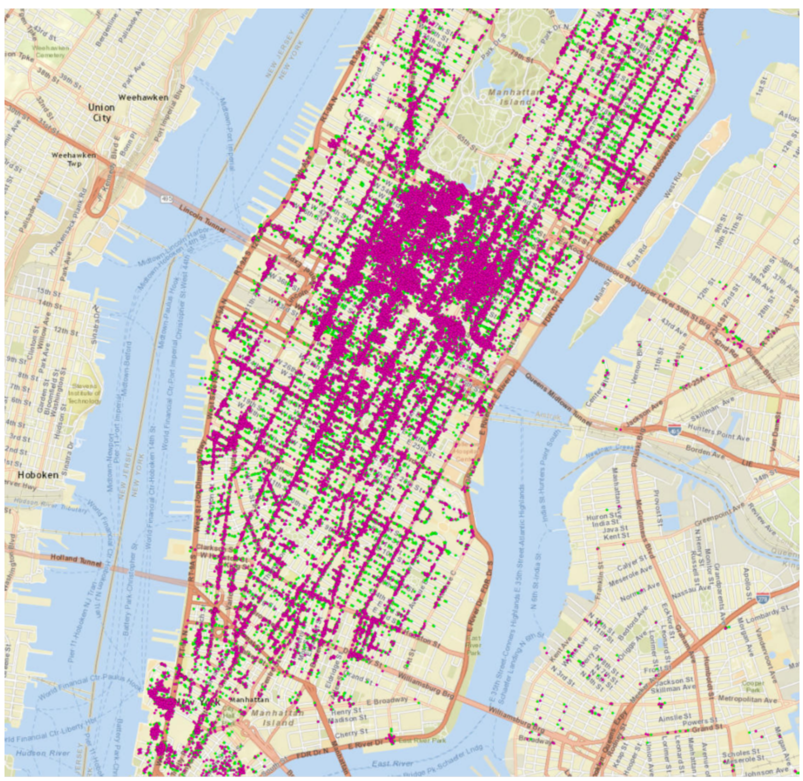

B Tables

Table 6: Robustness: Various Time Windows

	D.V. = Taxi dist./Uber dist.				
	5 min.	**10 min.**	**15 min.**	**20 min.**	**30 min.**
Airport	0.074***	0.073***	0.074***	0.073***	0.072***
	(0.005)	(0.003)	(0.003)	(0.003)	(0.002)
Non_local	-0.002	-0.003	0.001	-0.001	-0.003
	(0.005)	(0.003)	(0.003)	(0.002)	(0.002)
Airport × Non_local	0.015**	0.016***	0.013***	0.015***	0.016***
	(0.006)	(0.004)	(0.003)	(0.003)	(0.002)
Surge_multiplier	0.009	0.002	-0.000	0.001	0.002
	(0.008)	(0.005)	(0.004)	(0.003)	(0.003)
Surge_multiplier × Airport	-0.023**	-0.024***	-0.022***	-0.020***	-0.024***
	(0.011)	(0.007)	(0.005)	(0.005)	(0.004)
N	32,475	60,124	86,627	112,179	161,896
# of Uber drivers	15,565	20,390	23,109	24,770	27,028
R^2	0.562	0.440	0.376	0.336	0.285
	D.V. = Taxi dur./Uber dur.				
	5 min.	**10 min.**	**15 min.**	**20 min.**	**30 min.**
Airport	0.034***	0.032***	0.036***	0.038***	0.037***
	(0.008)	(0.005)	(0.003)	(0.003)	(0.003)
Non_local	0.000	-0.001	0.006	0.005	0.005*
	(0.008)	(0.005)	(0.004)	(0.004)	(0.003)
Airport × Non_local	-0.002	0.000	-0.008*	-0.009**	-0.011***
	(0.010)	(0.006)	(0.005)	(0.004)	(0.003)
Surge_multiplier	-0.006	-0.001	-0.002	-0.003	-0.005
	(0.013)	(0.007)	(0.006)	(0.005)	(0.004)
Surge_multiplier × Airport	-0.006	-0.003	-0.006	-0.001	-0.002
	(0.018)	(0.011)	(0.009)	(0.007)	(0.006)
N	32,475	60,124	86,627	112,179	161,896
# of Uber drivers	15,565	20,390	23,109	24,770	27,028
R^2	0.533	0.401	0.330	0.288	0.233

Notes. Controls not reported in the table include logged Uber driver total trips and logged Uber rider total trips. They are in expected signs and similar to the estimates reported in the main analyses. Hour-of-week fixed effects and Uber driver fixed effects are included in all specifications.
For all specifications, standard errors are cluster-robust at the hour-of-week level.
*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

Table 7: Robustness: Vendor 1

	D.V. = Taxi dist./Uber dist.			D.V. = Taxi dur./Uber dur.		
	Vendor1	**Vendor2 (Main sample)**	**Vendor1&2**	**Vendor1**	**Vendor2 (Main sample)**	**Vendor1&2**
Airport	0.093***	0.074***	0.082***	0.028***	0.036***	0.033***
	(0.003)	(0.003)	(0.002)	(0.004)	(0.003)	(0.003)
Non_local	0.003	0.001	0.000	0.005	0.006	0.005*
	(0.003)	(0.003)	(0.002)	(0.004)	(0.004)	(0.003)
Airport × Non_local	0.007**	0.013***	0.012***	-0.007	-0.008*	-0.008**
	(0.004)	(0.003)	(0.002)	(0.005)	(0.005)	(0.004)
Surge_multiplier	0.007	-0.000	0.001	-0.002	-0.002	0.000
	(0.005)	(0.004)	(0.003)	(0.007)	(0.006)	(0.004)
Surge_multiplier × Airport	-0.032***	-0.022***	-0.025***	-0.022**	-0.006	-0.014**
	(0.007)	(0.005)	(0.004)	(0.009)	(0.009)	(0.006)
Log (Uber_driver_total_trips)	0.007***	0.005**	0.005***	0.013***	0.012***	0.012***
	(0.002)	(0.002)	(0.001)	(0.003)	(0.003)	(0.002)
Log (Uber_rider_total_trips)	0.000	0.001	0.000	0.002***	0.003***	0.002***
	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)
N	71,850	86,627	158,477	71,850	86,627	158,477
R^2	0.413	0.376	0.296	0.362	0.330	0.239

Notes. Hour-of-week fixed effects and Uber driver fixed effects are included in all specifications.
For all specifications, standard errors are cluster-robust at the hour-of-week level.
*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

Table 8: Former Cabbies 2013

	D.V. = Taxi dist./Uber dist.			D.V. = Taxi dur./Uber dur.		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Airport	0.096***	0.105***	0.098***	0.008	-0.003	-0.029
	(0.011)	(0.013)	(0.021)	(0.019)	(0.020)	(0.034)
Non_local	-0.007	-0.009	-0.010	-0.020*	-0.023*	-0.021
	(0.007)	(0.007)	(0.013)	(0.011)	(0.012)	(0.019)
Airport × Non_local	0.022*	0.021	0.021	0.015	0.016	0.043
	(0.013)	(0.014)	(0.025)	(0.023)	(0.024)	(0.041)
Surge_multiplier	0.008	0.006	0.008	0.012	0.009	-0.002
	(0.009)	(0.009)	(0.017)	(0.013)	(0.014)	(0.025)
Surge_multiplier × Airport	-0.004	-0.000	-0.028	-0.072*	-0.076*	-0.023
	(0.022)	(0.021)	(0.033)	(0.041)	(0.042)	(0.068)
Log (Uber_driver_total_trips)	-0.003	-0.003	-0.003	0.001	-0.000	-0.001
	(0.002)	(0.002)	(0.004)	(0.004)	(0.004)	(0.006)
Log (Uber_rider_total_trips)	0.002	0.001	0.002	0.004	0.004	0.010*
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.006)
Hour of week FE	No	Yes	Yes	No	Yes	Yes
Uber driver FE	No	No	Yes	No	No	Yes
N	3,221	3,221	3,221	3,221	3,221	3,221
R^2	0.071	0.127	0.598	0.003	0.053	0.549

Notes. For all specifications, standard errors are cluster-robust at the hour-of-week level.
*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.

Table 9: Former Cabbies 2016

	D.V. = Taxi dist./Uber dist.			D.V. = Taxi dur./Uber dur.		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Airport	-0.004	-0.006	-0.006	0.016	0.019	0.036
	(0.007)	(0.008)	(0.012)	(0.012)	(0.014)	(0.023)
Non_local	0.004	0.002	-0.002	0.009	0.009	0.010
	(0.005)	(0.006)	(0.008)	(0.011)	(0.011)	(0.015)
Airport × Non_local	-0.001	0.003	0.009	-0.006	-0.005	-0.013
	(0.010)	(0.011)	(0.015)	(0.015)	(0.016)	(0.026)
Surge_multiplier	-0.001	0.001	-0.001	0.016	0.018	0.004
	(0.009)	(0.010)	(0.011)	(0.015)	(0.017)	(0.020)
Surge_multiplier × Airport	0.006	0.013	0.004	-0.049	-0.040	-0.041
	(0.026)	(0.025)	(0.032)	(0.032)	(0.036)	(0.054)
Log (Uber_driver_total_trips)	0.002	0.002	-0.004	0.008*	0.009**	0.007
	(0.003)	(0.003)	(0.010)	(0.004)	(0.005)	(0.018)
Uber_driver_rating	-0.037	-0.029	0.000	0.122***	0.150***	0.000
	(0.028)	(0.029)	(.)	(0.044)	(0.047)	(.)
Log (Uber_rider_total_trips)	-0.001	-0.002	-0.002	-0.003	-0.004	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)
Hour of week FE	No	Yes	Yes	No	Yes	Yes
Uber driver FE	No	No	Yes	No	No	Yes
N	3,751	3,751	3,751	3,751	3,751	3,751
R^2	0.001	0.056	0.484	0.004	0.051	0.511

Notes. For all specifications, standard errors are cluster-robust at the hour-of-week level.
*** Significant at the 1% level. ** Significant at the 5% level. * Significant at the 10% level.