Université Fédérale
Toulouse Midi-Pyrénées

THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par l'Université Toulouse 1 Capitole

Présentée et soutenue par

Emil PALIKOT

Le 5 juin 2020

Essays in Industrial Organization

Ecole doctorale : **Toulouse Sciences Economiques**

Spécialité : **Sciences Economiques - Toulouse**

Unité de recherche :
UMR 5314 -TSE-R

Thèse dirigée par
Marc IVALDI

Jury
M. Nathan Miller, Rapporteur
M. Paul Heidhues, Rapporteur
M. Bruno Jullien, Examinateur
M. Mathias Reynaert, Examinateur
M. Marc Ivaldi, Directeur de thèse

Essays in Industrial Organization

Ph.D Thesis

Emil Palikot*
Toulouse School of Economics

*e-mail address: emil.palikot@gmail.com.

Acknowledgements

This document is the final outcome of my graduate studies at Toulouse School of Economics. However, the six years spent in Toulouse gave me much more than the research projects composing this thesis. They gave me a marvelous journey, full of challenges and setbacks, but also marked with a few accomplishments; a period of intensive study accompanied by mentors and friends.

First and foremost, I would like to thank my mentors, an outstanding faculty of Toulouse School of Economics. Marc Ivaldi provided me with guidance and direction and supported me vigorously throughout the job market. I am deeply grateful for this. I have greatly benefitted from the unfading support and invaluable advice of Patrick Rey. I am also indebted to Jacques Crémer, Daniel Ershov, Bruno Jullien, Yassine Lefouili, and Mathias Reynaert, meetings and discussions with them shaped me as a researcher. They are my role models. I would like to thank Charles Angelucci and Timothy Simcoe for organizing visits to Columbia Business School and Questrom School of Business. These were truly enriching experiences, which reinforced my conviction to stay on the academic path.

Graduate studies are often described as a lonesome endeavor, that was not my experience. Each chapter of this thesis is co-authored with a fellow Ph.D. student. Collaborations with Matias Pietola, Xavier Lambin, and Rossi Abi Rafeh were the highlights of my graduate studies; they are excellent researchers and colleagues. In the pursuit of a Ph.D. I made friends who will forever stay close. Georgios, Rafa, Rossi, Matias, Tony, and Vicente you made Toulouse my favorite place. Constant encouragement from my brother, Aleksander, helped me stay focused and in good spirits. Finally, I would have never completed graduate studies without relentless support from my girlfriend Tota. Her kindness is unmatched.

Abstract

This thesis is composed of three chapters. The first chapter (co-authored with Matias Pietola) provides a theoretical model of pay-for-delay settlements with multiple firms. The second chapter (co-authored with Xavier Lambin) studies the impact of online reputation on ethnic discrimination. The third chapter (co-authored with Rossi Abi Rafeh) develops and estimates a model of industry dynamics.

The first chapter is motivated by recent antitrust cases in the pharmaceutical industry. It studies the interplay between pay-for-delay settlements, licensing deals and litigation. Our analysis highlights the externalities that they generate: pay-for-delay settlements reduce competition which encourages entry; licensing and litigation make entering less profitable. Faced with multiple entrants, the incumbent exploits these externalities by offering licensing deals to some entrants or by pursuing litigation in order to decrease the cost of delaying contracts offered to others. The number of delayed entrants increases with patent strength. Entrants without pay-for-delay settlements pursue litigation for patents of intermediate strength; otherwise, they receive licensing deals.

The second chapter shows that reputation systems can mitigate ethnic discrimination by enabling ethnic minority sellers to accrue a high reputation quickly, leading buyers to update their beliefs. Using data from a ridesharing platform, we find that minority drivers with no reviews make 12% less revenue relative to similar nonminority drivers. This disparity gradually shrinks and almost disappears for experienced drivers. To understand the mechanism behind this process, we construct a model of career concerns' of discriminated sellers in the presence of a reputation system. The model's estimates show that minority drivers, who just entered the platform, face overly pessimistic beliefs about the quality of their service. To alter these beliefs, they exert high effort and offer low introductory prices, swiftly boosting their reputation. Counterfactual simulations reveal that the cost of incorrect prior beliefs is high and that the reputation system strictly benefits minority drivers.

The final chapter studies the entry and pricing decisions of sellers in a market with a reputation system. We provide a model of dynamic oligopoly with heterogeneity in marginal and opportunity costs and individual reputation as a state variable. We show that new sellers are generally less likely to reenter the platform than incumbents and sellers who have a lower chance of entering in subsequent periods set on average higher prices. The mechanism behind these findings is selection on marginal costs. We apply our model to a dataset on sellers on a large ridesharing marketplace. We showcase a negative correlation of tenure on the platform, measured by the number of reviews, and prices set by drivers. However, after accounting for drivers' unobserved characteristics, which we interpret as marginal costs, we find a positive relationship. We provide, further, evidence of selection on unobservables by studying reentry decisions. Finally, we calibrate our dynamic model to uncover the distribution of opportunity costs.

Résumé

Cette thèse est composée de trois chapitres. Le premier chapitre (rédigé en collaboration avec Matias Pietola), fournit un modèle théorique de règlements 'pay-for-delay' avec plusieurs firmes. Le deuxième chapitre (en collaboration avec Xavier Lambin) étudie l'impact de la réputation digitale sur la discrimination ethnique. Le troisième chapitre (rédigé en collaboration avec Rossi Abi Rafeh) développe et estime un modèle de dynamique industrielle.

Le premier chapitre est motivé par les cas récents de politique de concurrence dans le domaine pharmaceutique et analyse les relations entre les règlements 'pay-for-delay', les contrats de licence et les litiges. Notre analyse met en lumière les externalitésde ces contrats: les règlements pay-for-delay réduisent la concurrence, ce qui encourage l'entrée dans le marché; les contrats de licence et le litige rendent l'entrée moins rentable. Avec plusieurs entrants en cours, le titulaire du brevet exploite ces externalités en proposant des accords de licence à certains entrants ou en poursuivant des contentieux afin de diminuer le coût des contrats dilatoires offerts aux autres. Le nombre des entrants tardifs augmente en fonction de la solidité du brevet. Les entrants sans règlements 'pay-for-delay' poursuivent des contentieux des brevets de solidité intermédiaire; sinon, ils reçoivent des accords de licence.

Le deuxième chapitre montre que les systèmes de réputation peuvent atténuer la discrimination ethnique en permettant les vendeurs appartenant à des minorités ethniques de construire une bonne réputation rapidement, ce qui entraîne les acheteurs d'actualiser leurs convictions. En utilisant une base de données collectée sur une plateforme de covoiturage, nous trouvons qu' en absence d'avis, les conducteurs membres des minorités ethniques gagnent 12% moins de revenue par rapport aux conducteurs non membres des minorités. Cette disparité diminue progressivement en fonction du nombre d'avis et disparaît presque complètement pour les conducteurs expérimentés. Pour comprendre le mécanisme derrière ce processus, nous concevons un modèle de 'career concerns' des vendeurs discriminés en présence d'un système de réputation. Les estimations du modèle montrent que les conducteurs appartenant à des minorités ethniques, qui viennent d'entrer dans la plateforme, font face à des convictions trop pessimistes quant à la qualité de leur service. Pour changer ces convictions, ils exercent de grands efforts et proposent des bas prix de lancement, pour renforcer rapidement leur réputation. Des simulations contrefactuelles révèlent que le coût des croyances antérieures erronées est élevé et que le système de réputation bénificie strictement aux conducteurs des minorités ethniques.

Le dernier chapitre étudie l'entrée sur un marché avec un système de réputation et les décisions de formation de prix des vendeurs. Nous proposons un modèle d'oligopole dynamique avec une hétérogénéité des coûts marginaux et des coûts d'opportunité, et avec la réputation individuelle comme une variable d'état. Nous montrons que les nouveaux vendeurs sont généralement moins susceptibles de rejoindre la plateforme par rapport aux anciens, et les vendeurs avec une faible chance de rejoindre dans les périodes suivantes mettent des prix moyens plus élevés. Le mécanisme derrière ces résultats est la sélection selon les coûts marginaux. Notre modèle s'appuie sur une base de données de vendeurs sur une grande plateforme d'un marché de covoiturage. Nous constatons une corrélation négative des utilisateurs expérimentés de la plateforme, mesurée par le nombre d'avis et les prix déterminés par les conducteurs. Cependant, après avoir pris en compte les caractéristiques non observées des conducteurs, dont lesquelles on interprète comme coûts marginaux, nous trouvons une relation positive. Par ailleurs, en étudiant les décisions d'une nouvelle entrée sur la plateforme, nous démontrons la sélection des non-observables. Finalement, nous découvrons la distribution des coûts d'opportunité, en calibrant notre modèle dynamique.

Contents

Chapter 1

Pay-for-Delay with Settlement Externalities

This chapter is based on a research project with Matias Pietola (Palikot and Pietola (2019)).

1.1 Introduction

A patent grants its owner the right to exclude potential competitors from the market for a limited period. When it expires, the market opens for entry and competition lowers prices to the benefit of consumers. By determining the scope and duration of patents, the intellectual property system aims to balance incentives to innovate against ex-post consumer surplus. In practice, however, patent offices have limited resources and grant questionable patents, which can be challenged in court.[1] As litigation is costly, the parties often settle out of court.

Patent settlements in the pharmaceutical industry have recently caught the attention of antitrust authorities around the world, with the so-called "pay-for-delay" agreements spurring a particularly heated debate. These are settlements between an incumbent patent holder and a potential entrant: in exchange for financial compensation, the entrant agrees not to challenge the validity of the patent and to stay out of the market for a certain period. Such deals fall at the intersection of antitrust and intellectual property policies. The European Commission considers pay-for-delay agreements to be anti-competitive and has imposed significant fines on companies involved, most notably in *Servier*[2] and *Lundbeck*.[3] The US Supreme Court has instead adopted a rule of reason approach, as in *Actavis*.[4]

However, patent disputes need not result in pay-for-delay settlements. Over the last few years, these agreements have constituted only around 3-12% of all patent settlements in the pharmaceutical sector in the European Economic Area.[5] Licensing agreements, where an entrant buys a license from the incumbent and enters without delay, are more frequent. Although pay-for-delay

[1] Farrell and Shapiro (2008) provide a detailed analysis of the economics of ''weak patents".

[2] European Commission decision, 9.7.2014, C(2014) 4955.

[3] European Commission decision, 19.6.2013, C(2013) 3803.

[4] *FTC v. Actavis, Inc.*, 133 S. Ct. 2223, 570 U.S. 136, 186 L. Ed. 2d 343 (2013).

[5] See The Pharmaceutical Sector Inquiry by the European Commission. From 2008, the European Commission has annually monitored patent settlements made by pharmaceutical companies in the EEA area.

agreements and licensing deals often coexist (with different entrants), the anticompetitive effects of the former have been analyzed in isolation. To our knowledge, the question of why incumbents offer different deals to (often similar) entrants has not yet been addressed in the economic literature. This article aims to fill the gap.

We develop a setting with one incumbent and multiple identical, potential entrants. The incumbent owns a patent of uncertain validity and enjoys a legal monopoly unless a court declares the patent invalid. Each entrant can either litigate, wait for the market to open (i.e., for the patent to expire or to be invalidated following litigation by another entrant), or settle with the incumbent. A settlement deal includes a financial transfer and an entry date.

Our analysis highlights the role of the externalities of settlement agreements. An entrant accepting a settlement with a late entry date imposes a positive externality on all other entrants: the expected profit attainable through litigation increases (due to reduced competition), which enhances the bargaining position of entrants negotiating licensing agreements or pay-for-delay deals. In particular, to monopolize the market (i.e., to delay the entry of all potential competitors), the incumbent would have to compensate every entrant for forgoing the duopoly profit that it could achieve through litigation. When there are many entrants, this cost exceeds the gain from maintaining the monopoly position.

Instead of delaying all entrants, to some of them the incumbent may license the patent or take them to court and face the risk of invalidation. Both strategies lower the cost of delaying other entrants, as they face increased competition if they reject their settlement offers and enter through litigation. Licensing and litigation are substitutes for the incumbent and never coincide. Litigation is costly but leads to higher profits if the court upholds the patent.

We show that the number of delayed entrants is higher when litigation costs are high, and the patent is strong. The incumbent may adopt a divide-and-conquer strategy, where it pays some entrants to delay their entry and either grants licenses to the others or fights them in court. Patents of intermediate strength are litigated, whereas sufficiently weak and strong patents are licensed. Furthermore, litigation is more likely to occur when litigation costs are low. In the extreme case, where litigation is costless, it is always the equilibrium outcome provided that competition is not too intense.

Our results contrast with predictions from single-entrant models. When facing only one entrant, the incumbent always delays entry; this strategy maximizes the industry profit (the market remains a monopoly), which is then shared between the parties according to their bargaining powers (Shapiro, 2003). However, with multiple entrants, the cost of offering pay-for-delay settlements to all of them may well exceed the gain from monopolization.

We provide an extension of our model where we allow the incumbent to offer settlements contingent on the validity of the patent. That is, an entrant agreeing to a pay-for-delay deal can nevertheless enter the market if the patent is invalidated, following litigation by another entrant. Therefore, by accepting a settlement, an entrant does not forgo the profit from entry in the case of patent invalidation. In addition, the payoff of a successful litigator decreases, as all competitors enter after the patent is invalidated. Hence, the payments required to delay entry are lower. As a result, the incumbent will delay all entrants in equilibrium, regardless of the strength of the patent and the costs of litigation. Thus, settlements contingent on patent validity hamper competition, rather than promote it.[6]

The possibility of entering into pay-for-delay agreements may also encourage the incumbent to license the patent, which benefits consumers. Because of this interdependency, the analysis of the welfare consequences of prohibiting pay-for-delay settlements should not look at them in isolation, but account for the decreased incentives to license the patent. To make this point, we provide numerical examples illustrating the possibility that a ban on pay-for-delay settlements may decrease consumer welfare.

Related literature The current economic literature on patent settlements has two branches, one for pay-for-delay agreements (Edlin et al., 2015; Elhauge

[6] In *Servier*, parties subject to a pay-for-delay agreement conditioned on patent validity argued that such a contract is pro-competitive. We show that the exact opposite is true.

and Krueger, 2012; Gratz, 2012; Manganelli, 2014; Meunier and Padilla, 2015; Shapiro, 2003) and another for licensing deals (Amir et al., 2014; Farrell and Shapiro, 2008; Kamien and Tauman, 1986; Katz and Shapiro, 1987; Lemley and Shapiro, 2005). To our knowledge, licensing and pay-for-delay agreements have not previously been studied together. In this article, we show that this obfuscates an important economic mechanism, triggered by the settlement externalities between entrants. Because of these externalities, the incumbent may offer different agreements to otherwise similar entrants. This observation relates our work to the literature on contracting with externalities (Segal, 1999, 2003).

Shapiro (2003) introduces the canonical model of pay-for-delay. He considers a framework with a single entrant who may challenge the incumbent patent holder. The two parties have the opportunity to settle and avoid going to court. Generally, they will conclude a pay-for-delay settlement which extends the monopoly period, and divide the resulting high profits. Our approach extends this seminal work by allowing for multiple entrants. This modification reveals the connection between pay-for-delay agreements, licensing, and litigation. In this way, we show that entry may occur in sequence, even though agreements are reached simultaneously with identical entrants.

Pay-for-delay settlements in an environment with multiple entrants have been previously studied by Meunier and Padilla (2015). They focus on the credibility of a litigation threat when a successful litigation by one firm opens the market for all entrants. Thus, entrants who do not pursue litigation effectively free-ride on the litigator. Our results also show the phenomenon of free-riding on litigation efforts, even though all entrants are symmetric and have the ability to litigate (Meunier and Padilla (2015) assume that only one of the entrants can start the litigation). Due to this logic, some patents are too strong to be challenged. We explicitly derive the threshold of patent strength above which entrants will not challenge the patent; however, our focus is on the cases where the threat of litigation is credible.

Shapiro (2003) proposes a general rule for evaluating patent settlements: allowing for settlements should not leave the consumers worse off compared to prohibiting them. Therefore, welfare analysis simplifies to comparing the duration of exclusion resulting from the settlement (the agreed entry date) to the expected entry date when there is no settlement (i.e., entry can occur through expiry or the invalidation of the patent). When the reverse payment is higher than the litigation cost incurred by the incumbent, exclusion due to settlement will exceed the expected delay from litigation (Shapiro, 2003). Elhauge and Krueger (2012) argue that all pay-for-delay settlements with a reverse payment higher than litigation costs should be illegal, regardless of the probability of the patent being invalid. We argue that this logic may fail when there is more than one entrant. We show that the fact of observing a high payment from the incumbent to an entrant does not necessarily imply that consumers are harmed.

Divide-and-conquer strategies have previously been studied in different contexts. Posner et al. (2010) show how a defendant can optimally exploit coordination failure between several plaintiffs. Other related works include Daughety and Reinganum (2002) and Che and Spier (2008). Typically, in these articles some of the plaintiffs are offered beneficial treatments and decide to settle with the defendant, which leads to the others dropping their lawsuits. In our case, the incumbent also exploits a coordination failure of the entrants, but the context and the modeling approach differs from the previous literature.

The article is organized as follows. The next section is devoted to case studies in the pharmaceutical industry. We discuss the context of pay-for-delay settlements and the antitrust response to them. Section 3 introduces the model. Section 4 presents an extension of our model, by allowing the incumbent to offer contracts contingent on the validity of the patent. Section 5 shows how a ban on pay-for-delay agreements changes the equilibrium outcome, and provides some policy implications, and Section 6 concludes. The Appendix contains all proofs.

1.2 Pay-for-delay cases in the pharmaceutical industry

The competitive landscape in the pharmaceutical sector is, to a large extent, shaped by two factors: the patent protection of newly developed drugs, and entry by producers of generic bio-equivalent medicines. Once the primary patent protecting the main chemical compound has expired, the producers of generics can simultaneously contest the originator's monopolistic position. From that date on, (usually) several generic producers race to be the first in the market. Often, despite the expiry of the main patent, the legal situation is, however, unclear: the originator has applied for secondary patent protection and sometimes holds patents on alternative methods of manufacturing the medicine. This legal uncertainty can result in a dispute over the validity of one of the patents that, in many cases, ends in a court of law. In this section, we briefly review three important cases involving pay-for-delay settlements; the first two are the European *Servier* and *Lundbeck* cases, and the last is a US case: *Actavis*.

The *Lundbeck* case considers several agreements between a Danish pharmaceutical company, Lundbeck, and several producers of generic drugs. In the 1970s and 1980s, Lundbeck developed the antidepressant drug Citalopram that was first marketed in the 1990s. The medicine was very successful and became the main product of Lundbeck; for example, it constituted 80-90% of the company's revenues in 2002. At the time of the settlements, during 2002 and 2003, patents related to the chemical compound and the original process had expired. In principle, the market was therefore free for generic producers to enter. However, Lundbeck still had a number of patents related to more efficient or alternative ways of manufacturing the drug. In order to limit the extent of market entry by generic producers, Lundbeck implemented a so-called "generic strategy" that involved different kinds of agreements with several entrants. Overall, the "generic strategy" was a mixture of pay-for-delay settlements, takeovers, licensing, accommodation, and even an introduction of own authorized generic producers. For example, in the UK, it allowed one firm to enter the market but offered a reverse payment to another. In Iceland, Lundbeck allowed a market entry without litigation. In June 2013, the European Commission ruled that the part of the "generic strategy" consisting of settlements delaying market entry was a violation of the European antitrust law, as its objective was to reduce potential competition. Lundbeck received a fine of EUR 93.8 million, and the generic producers involved were fined EUR 52.2 million.

The *Servier* case involves a French pharmaceutical manufacturer, Servier, and generic producers of Perindopril, a medicine for treating high blood pressure developed by Servier in the 1980s. Perindopril became Servier's most successful product with annual global sales exceeding USD 1 billion in 2006 and 2007, with average operating margins beyond 90%. Generic entry started to impose a credible threat to Servier once the patent governing the main compound expired in May 2003. Anticipating this, Servier started designing and implementing "the generic strategy" from the late 1990s. This strategy included acquiring new patents and resulted in five settlement agreements with different generic producers between 2005 and 2007. Four of these agreements were pay-for-delay settlements, whereas the fifth was a licensing deal. Servier considered litigation and licensing as alternative strategies to accommodate one of the challengers. In July 2014, the European Commission imposed fines totaling EUR 427 million on Servier and other pharmaceutical companies involved in the pay-for-delay settlements.

The Servier and Lundbeck decisions were both appealed before the EU General Court, which gave its judgment on the Servier case in December 2018. For the most part, the court upheld the Commission's decision, except for the Krka (a manufacturer of a generic version of Perindopril) settlement, for which the fine was annulled. The court made a clear distinction between side-deals in general and licensing agreements between Servier and Krka. Most interestingly, the court stated that licensing agreements should not be considered as suspicious side-deals, but as an appropriate means of settling a patent dispute: a licensing agreement stems from the parties' mutual understanding about the validity of the patent, making the settlement possible.[7]

In the US, the Drug Price Competition and Patent Term Restoration Act of 1984 (the so-called Hatch-Waxman Act) shapes the regulatory approval of generic drugs. This legislation aims to promote the entry of generics by guaranteeing the first of them a duopolistic position. When a producer of generics files an Abbreviated New Drug Application (ANDA) to the Food and Drug Administration (FDA), the Hatch-Waxman Act requires the declaration of a relationship to a patent mentioned in the Orange Book.[8] If a generic producer states that the relevant patents are no longer valid, or that they are not infringed, the certification is granted. Importantly, the Hatch-Waxman Act provides 180 days of exclusivity to the first entrant: no other producer of generics can obtain approval from the FDA during this time.

The FTC vs. Actavis case was brought to the US Supreme Court by the Federal Trade Commission (FTC) in 2013. The case considers a deal made between a Belgian pharmaceutical company, Solvay Pharmaceuticals, and a generic producer, Actavis Inc. Solvay was granted a new patent for AndroGel in 2003. Later, Actavis filed an ANDA to the FDA and stated that Solvay's new patent was invalid and thus the generic version produced by Actavis could not infringe upon the AndroGel patent. Solvay settled the case with Actavis: the settlement agreement included a reverse payment from Solvay to Actavis in return for an exclusion period during which Actavis agreed to stay out of the market. The agreed entry date was 65 months before the AndroGel patent expired. The FTC considered the arrangement between Solvay and Actavis as an antitrust violation and brought a lawsuit against them. The District Court and the appellate court, the Eleventh Circuit, dismissed the case. However, the US Supreme Court overturned their decisions and held that it was not sufficient to base the legal analysis on patent law policy and that the antitrust question must be addressed. The Supreme Court argued that: (1) FTC's complaint could not have been dismissed without analyzing the potential justifications

[7]Case T-691/14, Servier v. Commission, EU:T:2018:922, paras 943-963.

[8]A list of approved drug products together with a catalog of patents related to each of them.

for such a decision; (2) the patent holder was likely to have enough power to implement antitrust harm in practice; (3) the antitrust action was likely to prove more feasible administratively than the Eleventh Circuit believed: a large, unexplained payment from the patent holder to the generic producer could provide a workable surrogate for a patent's weakness; and (4) the parties could have made another type of settlement agreement by allowing the generic producer to enter the market before the patent expired, without the need for a reverse payment.

Several valuable lessons can be drawn from the case law on pay-for-delay agreements. First of all, even a weak patent can be useful for the incumbent, because of high litigation costs and free-riding between the generic companies in their litigation efforts. Patent invalidation opens the market for everybody, not merely for the generic producer who took the litigation effort and, incurred an often significant litigation cost. The generic entrants have expressed their concern to "win the battle, but lose the war" due to follow-up entry to the market.[9] The incentives to settle litigation are particularly pronounced when other generic producers wait to enter the market.

Second, the terms of a settlement have to reflect both the competitive situation in the market and the strength of the patent. In order to reach a mutually beneficial settlement, parties need to have a similar assessment of the strength of the patent. To ensure that these assessments reflect the actual probability of patent invalidation, parties undertake laboratory tests and seek third-party advice.[10]

Third, the entry game starts after the expiry of a certain patent, and this date is common knowledge. Firms, wishing to enter the market arrive simultaneously, and their subsequent sequential entry is an outcome of an interplay between the patent holder and the entrants. Some entrants may receive a license or go to court, whereas the others are delayed. So-called "generic strategies" are typically combinations of pay-for-delay agreements, licensing deals, litigation, and takeovers.

[9] See paragraph 493 of the Servier decision by the European Commission.

[10] See paragraph 709 of the Servier decision and 522 of the Lundbeck decision.

1.3 Model

Consider a market with one incumbent firm, I, and $N \geq 2$ symmetric entrants, E_1, E_2, ..., E_N. I owns a patent and enjoys a legal monopoly until the patent expires, unless at least one E_i, for $i \in \{1, 2, \ldots, N\}$, litigates and a court declares the patent invalid. Litigation costs $c \geq 0$ to E_i and $C \geq 0$ to I. It is common knowledge that if any E_i litigates, the court will declare the patent invalid with probability $1 - \theta$ and uphold the patent otherwise; θ thus reflects the strength of the patent. We assume that, if several entrants litigate, the litigation outcomes are perfectly correlated.[11]

The patent starts on date $t = 0$, and time runs continuously until the patent expires at time $t = 1$, after which free entry drives all profits down to zero. At date zero, the firms play a two-stage negotiation game:

- **Stage 1**: I offers each E_i a settlement deal, which includes an entry date $t_i \in [0, 1]$ and a payment $p_i \in \mathbb{R}$ from E_i to I. The offers are public and are observed by all firms.

- **Stage 2**: Each E_i either accepts the settlement offer, rejects the offer and litigates, or rejects the offer and waits for the market to open.

Litigation is time-consuming: it takes the court $1 - l$ time to make a decision; thus, after the court rules on patent validity there is $l > 0$ time left before the initial patent expiry date. If the court declares the patent invalid, those entrants who have rejected the settlement offer (either litigate or wait) enter the market, whereas those who have signed a settlement deal, are bound by it.[12] If instead the court upholds the patent, entry is possible only through a settlement. This is also the case during the litigation period.[13]

[11]In practice, courts often bundle similar cases.

[12]This is in line with the legal principle of *pacta sunt servanda*. It is the ex-ante view of the strength of the patent that matters for reaching a settlement before the court, not the validity of the patent resolved ex-post. Later on we consider entry delay contingent on patent validity.

[13]We thus do not allow for entry at risk, which would open the question about the appropriate damage rule if the patent is valid.

At any point in time, if n firms are active in the market (I and $n-1$ entrants), I makes a profit $\Pi(n)$ and each active E_i makes a profit $\pi(n)$, whereas the entrants that stay out make zero profit.[14] We assume that the profits are all positive and that total industry profit decreases in n. All firms are risk-neutral, and their payoffs are the sums of payments and profits. The equilibrium concept is subgame perfect Nash equilibrium in pure strategies and we use backward induction to solve the game.

We make two key assumptions. First, we assume that the monopoly profit is not too high:

Assumption 1. *There exists $x \in \{0, \ldots, N-1\}$ such that*

$$\Pi(1+N-x) - x\pi(2+N-x) > \Pi(1) - N\pi(2).$$

The purpose of this assumption is to rule out the trivial case when I always finds it profitable to delay all entry until patent expiration, even if the patent is invalid with certainty and litigation is costless. The assumption is satisfied by the symmetric Cournot quantity-setting game with a linear demand, for example, and generally satisfied for N large enough.

Second, we assume that the litigation threat is credible:

Assumption 2. *The expected payoff from litigation is at least zero:*

$$(1-\theta)\, l\pi(1+N) - c \geq 0.$$

This assumption is necessary to make the game interesting, as I would otherwise simply make unacceptable settlement offers to all entrants and monopolize the market at no cost.

Although, I can offer settlements with any entry date $t_i \in [0,1]$ and payment $p_i \in \mathbb{R}$, only two types of settlement agreements are signed in equilibrium (see the Appendix for the formal proof):

[14] Note that the incumbent may have a different profit than the entrants, capturing entry costs and product differentiation, for example between a branded drug and a generic drug in the pharmaceutical industry.

- Licensing agreements: E_i enters at date $t_i = 1 - l$ and pays a licensing fee $p_i > 0$;
- Pay-for-delay agreements: E_i delays entry until patent expiry $t_i = 1$ and receives a reverse payment $-p_i > 0$.

There are two reasons for this. First, entry before the court decides on the validity of the patent is not possible without a settlement. Thus, as the industry profits are highest under monopoly, I never finds it profitable to license during the litigation period. Second, we do not allow for contracts with exit, implying that payoffs are time-independent. Hence, if licensing is profitable for I at some point, it is already profitable at the moment when the litigation period ends.

To characterize the equilibrium of the negotiation game, it is useful to define two functions: $X : [0,1] \longrightarrow \{0, 1, 2, \ldots, N\}$ and $f : [0,1] \longrightarrow \mathbb{R}$, where $X(\theta)$ is a function of the strength of the patent, and it maximizes:

$$\begin{aligned} g(x;\theta) \coloneqq\; & \theta \left[(N - x)\, \pi (1 + N - x) + x\pi (2 + N - x) \right] \\ & + \Pi (1 + N - x) - x\pi (2 + N - x) . \end{aligned} \tag{1.1}$$

As it turns out, $X(\theta)$ determines the number of delayed entrants in equilibrium.

$$f(\theta) \coloneqq \theta \Pi (1) + (1 - \theta)\, g(X(0); 0) - g(X(\theta); \theta), \tag{1.2}$$

is the difference in maximum profits I can attain by pursuing litigation and by settling with all entrants, net of total effective litigation costs λ, where:

$$\lambda \coloneqq \frac{C + Nc}{l}. \tag{1.3}$$

$f(\theta) = \lambda$ makes I indifferent between N settlements and litigation. We are now ready to state our first result, which characterizes the equilibrium of the negotiation game, fixing the strength of the patent, and varying λ.

Proposition 1. *There exists an essentially unique equilibrium of the negotiation game:*

- *If litigation costs are high, $\lambda > f(\theta)$, then there is no litigation: $X(\theta)$ entrants are delayed and $N - X(\theta)$ buy a license;*

- *If litigation costs are low, $\lambda < f(\theta)$, then there is litigation: one entrant litigates, $X(0)$ entrants are delayed and $N - 1 - X(0)$ wait.*

Furthermore, $X(\cdot)$ is weakly increasing and satisfies $X(1) = N$.

As the litigation outcomes are perfectly correlated, at most one entrant litigates in equilibrium (due to symmetry, the identity of the entrant does not matter for I). Furthermore, as the litigation threat is credible, waiting is never a best response if no rival entrant litigates. Thus, in any equilibrium, either one and only one entrant litigates, or all entrants settle.

Whether litigation happens in the equilibrium depends on how high the litigation costs are relative to profits. For high litigation costs, $\lambda > f(\theta)$, all entrants settle. To delay x entrants, I must compensate each of them for payoff they would obtain by rejecting the settlement deal and going to court instead. The cost of entry delay is therefore:

$$x \cdot \underbrace{[(1-\theta)\, l\pi\,(2+N-x) - c]}_{\text{reverse payment}}. \tag{1.4}$$

Furthermore, the highest possible licensing fee each licensee is willing to pay corresponds to the difference between the profit they make in the market and their payoff from litigation. Thus, the licensing revenue is:

$$(N - x) \cdot \underbrace{[\theta l\pi\,(1+N-x) + c]}_{\text{licensing fee}}. \tag{1.5}$$

Reverse payments and licensing fees are increasing in the number of delayed entrants. Each entrant that accepts a delaying settlement imposes a positive externality on all other entrants: by agreeing to stay out of the market, a delayed entrant increases the expected profits from litigation to all other entrants, which in consequence increases settlement payoffs needed to compensate for withdrawing from litigation. Licensing fees are increasing in the number of delayed entrants, because profits in the market are higher when more entrants are delayed. Licensing agreements have an opposite effect: each entrant that accepts a licensing agreement imposes a negative externality on all other entrants; there is increased competition following a successful litigation. By accepting a settlement and forgoing litigation each entrants saves on the litigation cost; however, I accounts for it in settlement payments; thus, effectively I extracts all savings from forgone litigation costs.

In addition to the payments from the settlements, from the market I makes the profit:

$$\underbrace{(1-l)\,\Pi\,(1)}_{\text{litigation period}} + \underbrace{l\Pi\,(1+N-x)}_{\text{own profit under licensing}}. \tag{1.6}$$

Hence, putting everything together, I's payoff from delaying x entrants and licensing the patent to the rest writes:

$$(1-l)\,\Pi\,(1) + lg\,(x;\theta) + Nc, \tag{1.7}$$

which is maximized at $X\,(\theta)$.

The number of delayed entrants is weakly increasing in the strength of the patent. First, the cost of entry delay is decreasing in the strength of the patent. I has to compensate delayed entrants for withdrawing from litigation; the stronger the patent, the lower the expected profit from starting a litigation, and hence, the lower the necessary reverse payments. Second, industry profits are decreasing in the number of firms in the market. Thus, if the cost of entry delay decreases sufficiently, I always has an incentive to delay another entrant.

For low litigation costs, $\lambda < f\,(\theta)$, there is litigation in the equilibrium. In this case, to delay x entrants, I must compensate each of them for the payoff they would obtain by waiting and free-riding on the litigation costs taken by the entrant who went to court. The cost of entry delay writes:

$$x \cdot \underbrace{[(1-\theta)\,l\pi\,(2+N-x)]}_{\text{reverse payment}}. \tag{1.8}$$

In principle, entrants who are not delayed could either obtain a license or wait for the market to open through litigation. However, when litigation is already ongoing, I is always better off by making no deal than licensing the patent. To see this, suppose I offers a licensing deal to $y \leq N - x - 1$ entrants, yielding the licensing revenue:

$$y \cdot \underbrace{[\theta l \pi (1+y)]}_{\text{licensing fee}}. \tag{1.9}$$

Furthermore, from the market I would make the expected profit:

$$\underbrace{(1-l)\,\Pi(1)}_{\text{litigation period}} + \theta \cdot \underbrace{l\Pi(1+y)}_{\text{own profit if wins in court}} + (1-\theta) \cdot \underbrace{l\Pi(1+N-x)}_{\text{own profit if loses in court}}. \tag{1.10}$$

The number of licensees only affects the licensing revenue and I's own profit. As the industry profit is decreasing in the number of firms, the optimal number of licensees is zero: $y = 0$, which allows I to maintain a monopoly position if the court upholds the patent.

Generally, I takes entrants to court or licenses the patent in order to decrease the cost of entry delay. The cost of delaying entry equals compensation to all delayed entrants for withdrawing from litigation. Litigation brings them profits only if the court invalidates the patent. Therefore, the fact that licensing leads to entry for sure (unlike litigation) does not reduce the cost of delay, but decreases I profit for sure, rather than with probability $1 - \theta$. Hence, if litigation is already underway, I will not offer any licensing contracts.

However, when there is litigation I may still find it profitable to delay some of the entrants. Considering the payments and profits together, I's payoff from litigation is:

$$(1 - l + \theta l)\,\Pi(1) + (1 - \theta)\, l g(x; 0) - C, \tag{1.11}$$

which is maximized at $X(0)$.[15]

Finally, the difference between the highest profits I can obtain with and

[15]Note that $X(0) \leq N - 1$ by Assumption 1.

without litigation, net of the effective litigation costs, is:

$$f(\theta) = \theta\Pi(1) + (1-\theta)\, g(X(0); 0) - g(X(\theta); \theta). \tag{1.12}$$

Therefore, the difference between $f(\theta)$ and costs λ determines the equilibrium. If litigation is costly, $\lambda > f(\theta)$, the settlement payoff exceeds the payoff from litigation, and if the reverse inequality holds, I prefers to litigate.

By observing that $g(X(\theta); \theta)$ is convex, being the upper envelope of affine functions, we deduce that $f(\theta)$ is concave and continuous. Furthermore, it clearly satisfies $f(0) = f(1) = 0$. Hence, we obtain our next result, which characterizes the equilibrium of the game, fixing the costs of litigation, and varying θ.

Proposition 2. *For any $\lambda < \max f$ there exist thresholds of patent strength, $\underline{\theta}(\lambda) < \overline{\theta}(\lambda)$, where $\underline{\theta}(\lambda)$ satisfies $\underline{\theta}(0) = 0$ and increases in λ, and $\overline{\theta}(\lambda)$ satisfies $\overline{\theta}(0) = 1$ and decreases in λ, such that, in equilibrium:*

- *Patents of intermediate strength, $\theta \in \left[\underline{\theta}(\lambda), \overline{\theta}(\lambda)\right]$, are litigated;*
- *Patents that are sufficiently strong, $\theta > \overline{\theta}(\lambda)$, or weak, $\theta < \underline{\theta}(\lambda)$, are not taken to court: all entrants settle.*

Going to court is costly but gives I a chance of monopolizing the market without paying the entrants. If the patent is strong the entrants, who are likely to lose in court are willing to accept pay-for-delay agreements with small reverse payments. I then prefers to avoid costly litigation and to delay entrants. If instead the patent is weak, I's chance of monopolizing the market through litigation is small, and the entrants have a strong bargaining position. Therefore, for a weak enough patent, I offers licensing deals to save on the costs of litigation.

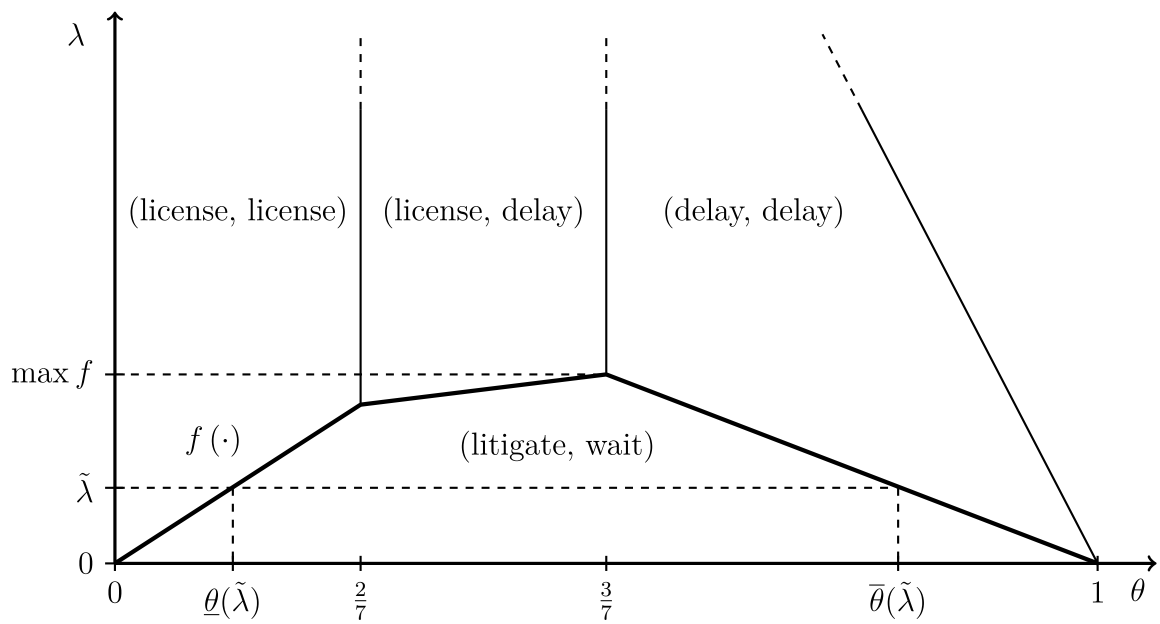

Figure 1.1: Equilibrium of the negotiation game with two entrants, as a function of the strength of the patent, θ, and the total cost of litigation, λ. Profits as in the symmetric Cournot quantity-setting game with an inverse demand $p = 1 - Q$. Details of this example are in the Appendix.

For patents of intermediate strength, I has a real chance of monopolizing the market through litigation, whereas delaying entry is not cheap. It will then take its chances and pursue litigation, unless the litigation costs are too high, in which case it is better to play "divide-and-conquer" by offering a licensing deal to some of the entrants while delaying the other ones.

Figure 1.1 uses a parametrized Cournot quantity setting game with two entrants to depict the outcome of the negotiation game as a function of the strength of the patent and the costs of litigation. For any level of costs λ we can determine the interval $\left[\underline{\theta}(\lambda), \overline{\theta}(\lambda)\right]$ of patent strength, where there is litigation in the equilibrium. For zero litigation costs, going to court is always the profit-maximizing strategy for the incumbent, and for sufficiently high litigation costs, there is never litigation.

Remark: Identical entrants have different equilibrium payoffs when I treats them differently. If there is litigation in equilibrium, the entrants who wait obtain the highest payoff, $(1-\theta)\, l\pi\, (1+N-X\,(0))$, whereas the litigator gets the same but has to pay for the litigation cost. The delayed entrants get reverse payments equal to $(1-\theta)\, l\pi\, (2+N-X\,(0))$, because I has to compensate them for the payoff they *would* obtain by rejecting the settlement deal. Depending on how costly litigation is, the delayed entrants are better off than the litigator or vice versa. If instead there is no litigation in equilibrium, each delayed entrant obtains $(1-\theta)\, l\pi\, (2+N-X\,(\theta))-c$, whereas licensees receive $(1-\theta)\, l\pi\, (1+N-X\,(\theta))-c$. Thus, licensees are always better off than delayed entrants.

1.4 Conditional settlements

The legal principle of *pacta sunt servanda* states that a contract should stay in power despite an *expected* change of environment. Following this principle, our analysis to this point has assumed that a settlement stays in force even when a court declares the patent invalid. Indeed, when parties negotiate over a settlement to a patent dispute, they have a belief about the strength of the patent, and this influences the terms of the settlement. If the patent is later declared invalid by a court, following litigation by a third party, the parties to the contract have been aware of this risk when agreeing to the settlement. It is the ex-ante view of the strength of the patent that matters in reaching a settlement, not the validity of the patent resolved ex-post.

Even so, the parties could explicitly formulate a settlement agreement conditional on the validity of the patent. Allowing for such deals completely changes the equilibrium of the game:

Proposition 3. *The negotiation game with conditional settlement terms has a unique equilibrium, where all entrants are delayed until patent expiration, regardless of the strength of the patent and the costs of litigation.*

Conditioning pay-for-delay settlements on the validity of the patent decreases the payoff from litigation. If the court invalidates the patent, the litigating entrant has to compete with all delayed entrants; thus the incentives to start the litigation are lower. As a result, the cost of entry delay reduces to:

$$x \cdot \underbrace{[(1-\theta)\, l\pi\,(1+N) - c]}_{\text{reverse payment (conditional deals)}} \leq x \cdot \underbrace{[(1-\theta)\, l\pi\,(1+N-x) - c]}_{\text{reverse payment (regular deals)}} . \tag{1.13}$$

We have argued that the key mechanism forcing I to allow for some entry is the positive externality imposed by an entrant accepting a delaying settlement on other entrants' payoffs from litigation. Pay-for-delay settlements conditioned on the validity of the patent do not create such an externality, because they are no longer binding when a successful litigator enters the market.

In *Servier*, for example, firms argued that such conditional pay-for-delay agreements are less harmful to consumers because the delayed entrant can enter the market if the patent is declared invalid. This logic, however, is misleading. By using conditional settlements, the incumbent is able to reduce the cost of entry delay to the extent that it is always profitable to delay all entrants until patent expiration. We thus have a clear policy implication: conditional settlement terms hamper competition by reducing the number of firms in the market, and thus from an antitrust policy perspective, they should not be allowed.

1.5 Banning pay-for-delay agreements

In this section, we look at the implications of a ban on pay-for-delay agreements on the equilibrium of the negotiation game. One could argue that from the perspective of the patent system, the outcome, as illustrated by Figure 1.1, is reassuring. First, strong patents are not challenged in court and result in a monopoly, which is the point of issuing patents: a patent grants its owner the right to exclude rivals from the market, thus, it would be a waste of resources if firms were litigating over the validity of *ex-ante* strong patents.

Second, weak patents are licensed; therefore, they do not prevent entry to the market, which benefits consumers. Third, patents of intermediate strength are litigated, correcting the legal uncertainty created by the imperfect screening of patents.

A ban on pay-for-delay agreements alters this picture significantly. To see how the equilibrium of the negotiation game changes, define:

$$f_0(\theta) := \theta\left[\Pi(1) - \Pi(1+N) - N\pi(1+N)\right], \tag{1.14}$$

which is simply the function $f(\theta)$ under the constraint $X = 0$.

Proposition 4. *Suppose pay-for-delay agreements are banned. Then, there exists an essentially unique equilibrium of the game:*

- *If litigation is costly, $\lambda > f_0(\theta)$, then there is no litigation: all entrants buy a license;*
- *If litigation is not costly, $\lambda < f_0(\theta)$, then there is litigation: one entrant litigates while the others wait.*

Furthermore, for any $\lambda < f_0(1)$, there exists a threshold of patent strength, $\theta_0(\lambda)$, such that there is litigation if the patent is sufficiently strong, $\theta > \theta_0(\lambda)$, and no litigation otherwise.

When pay-for-delay agreements are not allowed, the only way I can avoid litigation is to license the patent to all entrants. As a consequence, the space for litigation increases, because this is the only way for the incumbent to monopolize the market. For strong patents, signing pay-for-delay agreements would save on the costs of litigation, while still resulting in the same market outcome with a high probability. Furthermore, the scope of licensing increases, because offering all entrants a license is the only way the incumbent can avoid litigation, if litigation is too expensive. This means that even strong patents are licensed, when litigation costs are high. Figure 1.2 illustrates the outcome of the negotiation game when pay-for-delay agreements are illegal and thus outside the incumbent's toolkit.

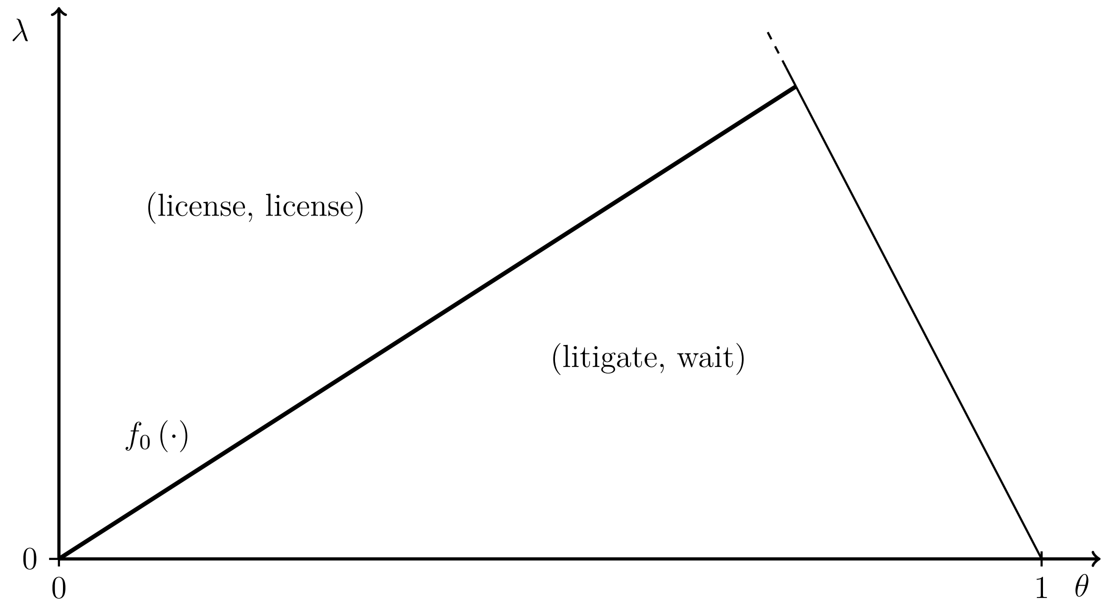

Figure 1.2: The equilibrium of the negotiation game with two entrants when pay-for-delay agreements are banned, as a function of the strength of the patent, θ, and the total cost of litigation, λ.

Consumer welfare in a market governed by a patent depends on the date of entry: licensing agreements result in earlier entry compared to litigation, and litigation leads to (expected) more competition than pay-for-delay. As we argue in this article, the possibility of entering into pay-for-delay agreements may encourage the incumbent to license the patent. Therefore, in some cases consumers might be worse-off from a ban on pay-for-delay settlements; this is when the outcome under the *laissez faire* rule involves licensing, and under the ban no licensing happens in the equilibrium. However, if all entrants are delayed, prohibiting pay-for-delay settlements would improve consumer welfare. Figure 1.3 illustrates how a ban on pay-for-delay agreements influences the consumer surplus ex-post in our parametrized example.

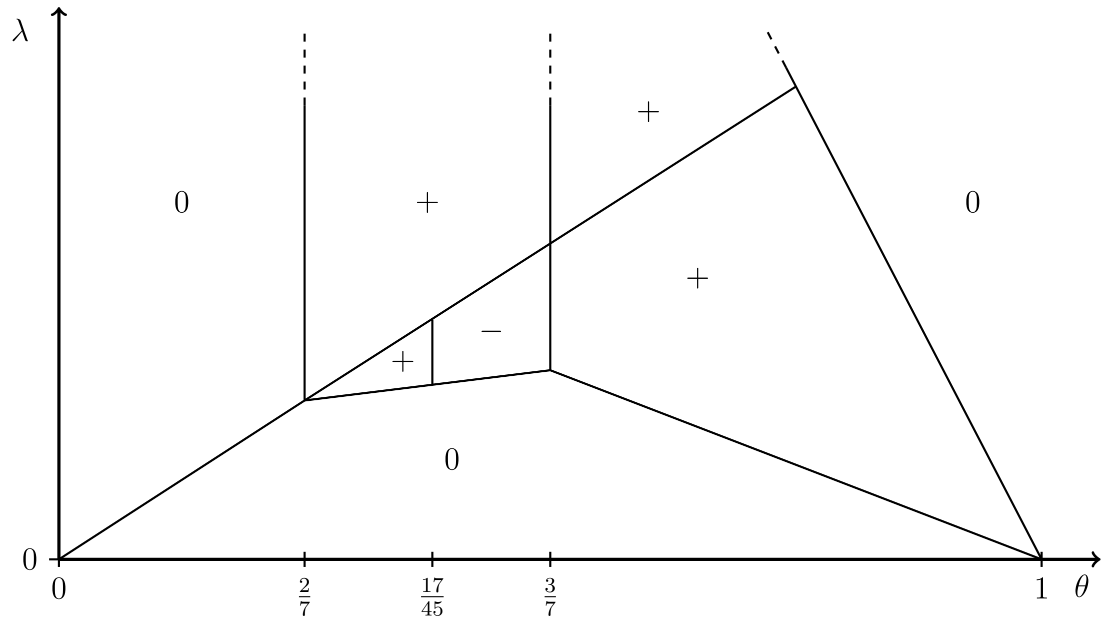

Figure 1.3: The change in expected consumer surplus (ex-post, not considering the incentives to innovate) following a ban on pay-for-delay agreements.

In the area with the negative sign the incumbent stops licensing the patent and goes to court instead, although the expected consumer surplus from litigation is lower than the one in duopoly, which would prevail if pay-for-delay agreements were allowed. This highlights a more general point: due to the interdependency between pay-for-delay and licensing incentives, the analysis of the welfare consequences of prohibiting pay-for-delay settlements should not look at them in isolation, but account for the decreased incentives to license the patent.

1.6 Concluding remarks

In a market covered by a patent, consumer welfare depends largely on the date of entry of competing firms. Therefore, settlements of patent litigation aimed at delaying or accelerating entry have a severe economic impact. We propose a model of patent settlements in an environment with multiple entrants. This approach allows us to study an important phenomenon which is new to the literature on pay-for-delay: settlement externalities. A settlement defining the date of entry of an entrant impacts incentives to enter the market by other firms. If a settlement leads to a delayed entry it incentivizes all other potential entrants to challenge the patent holder and enter the market; on the contrary, if a settlement is a licensing agreement, it discourages entry by others.

Economic literature studying pay-for-delay agreements has to date, focused either on the entry of a single firm (Shapiro, 2003) or on sequential entry (Gratz, 2012). It has been shown that an incumbent patent holder can agree with a potential entrant to delay its entry with a reverse payment (Shapiro, 2003). Allowing for multiple *ex-ante* identical entrants highlights an important externality of such an agreement, as it makes entering the market more attractive to all other potential entrants. In order to preserve its monopoly position, the incumbent has to offer expected duopoly profits to all potential entrants, as these are attainable by rejecting the settlement offer and pursuing litigation instead. For sufficiently many firms the patent holder will allow some entry.

We show that more pay-for-delay settlements are concluded when the patent is strong and the litigation costs are high. When delaying all entrants becomes too expensive, the incumbent will implement a more complicated strategy: divide-and-conquer. To decrease the cost of pay-for-delay settlements, the incumbent will allow some entry to the market, either through licensing or litigation. Thus, in the equilibrium, entrants receive different payoffs despite being identical. Furthermore, we show that sequential entry to the market can be an outcome of a negotiation game between an incumbent and several entrants arriving simultaneously. We, also, find that litigation occurs for patents of intermediate strength.

This article contributes to the debate on the economic consequences of pay-for-delay agreements. First, we prove that settlements which are conditional on patent validity should not be allowed if a policy-maker's goal is to promote entry to the market. When a settlement is only binding as long as the patent is valid, the positive externality caused by pay-for-delay settlements to the payoff from litigation disappears. Under conditional settlements, a firm entering the market through litigation faces competition from all entrants who have accepted delaying settlements, and achieves lower profits. As a consequence, reverse payments associated with pay-for-delay agreements are much lower, and the incumbent delays all entrants.

Second, by analyzing how the equilibrium of the game changes when pay-for-delay settlements are prohibited, we contribute to the discussion on the welfare consequences of pay-for-delay settlements. Previous literature studying models with one entrant provides useful, information-light rules to guide antitrust enforcement with respect to pay-for-delay settlements (Elhauge and Krueger, 2012; Shapiro, 2003). Unfortunately, in an environment with multiple entrants, the settlement date of entry or the size of payment are not sufficient statistics to determine the welfare consequences of concluded settlements. The interdependency of settlements due to the externalities discussed in this article means that such threshold rules might deem welfare-improving conduct as anti-competitive. In our setting, the *laissez-faire* approach to pay-for-delay settlements might improve consumer welfare when sufficiently many firms receive licenses.

A particular feature of the pharmaceutical market in the US is the Hatch-Waxman Act. This legislation aims to promote the entry of generics by guaranteeing the first entrant a duopoly position. In light of our results, such a policy should not be effective. Once exclusivity to the first entrant is granted, the incumbent will delay its entry; there are no settlement externalities because the other entrants are excluded from the market by law. However, a comprehensive study of such legislation should account for incentives to innovate by generics, which was outside the scope of our analysis.

We have explicitly avoided introducing asymmetric information. Beliefs about the strength of the patent are critical factors in agreeing on a settlement. Settlement offers in a game of incomplete information signal the strength of the patent, and the incumbent would have to account for the impact of offers on the beliefs. In a setting with multiple entrants, this could lead to complex settlement strategies, thus, careful analysis of such a signaling game is an interesting avenue for future research.

Finally, another valuable extension would be to introduce entry-at-risk. For example, in *Servier*, one of the entrants decided to launch its product before the resolution of the patent dispute. The incentive to "enter at risk" would be shaped by the possibility of obtaining an injunction and by the damage rule applied by the court. Determining conditions under which entry-at-risk could occur, and its welfare consequences is an important extension to the current analysis.

Bibliography

AMIR, R., D. ENCAOUA, AND Y. LEFOUILI (2014): "Optimal licensing of uncertain patents in the shadow of litigation," *Games and Economic Behavior*, 88, 320–338.

CHE, Y.-K. AND K. E. SPIER (2008): "Exploiting plaintiffs through settlement: Divide and conquer," *Journal of Institutional and Theoretical Economics JITE*, 164, 4–23.

DAUGHETY, A. F. AND J. F. REINGANUM (2002): "Informational externalities in settlement bargaining: confidentiality and correlated culpability," *RAND Journal of Economics*, 587–604.

EDLIN, A., S. HEMPHILL, H. HOVENKAMP, AND C. SHAPIRO (2015): "The Actavis Inference: Theory and Practice," *Rutgers University Law Review*, 67, 585–635.

ELHAUGE, E. AND A. KRUEGER (2012): "Solving the Patent Settlement Puzzle," *Texas Law Review*, 91, 284–330.

FARRELL, J. AND C. SHAPIRO (2008): "How Strong Are Weak Patents?" *American Economic Review*, 98, 1347–69.

GRATZ, L. (2012): "Economic Analysis of Pay-for-delay Settlements and Their Legal Ruling," Discussion papers in economics, University of Munich, Department of Economics.

KAMIEN, M. I. AND Y. TAUMAN (1986): "Fees Versus Royalties and the Private Value of a Patent," *The Quarterly Journal of Economics*, 101, 471–491.

KATZ, M. AND C. SHAPIRO (1987): "R&D Rivalry with Licensing or Imitation," *American Economic Review*, 77, 402–20.

LEMLEY, M. A. AND C. SHAPIRO (2005): "Probabilistic Patents," *Journal of Economic Perspectives*, 19, 75–98.

Manganelli, A.-G. (2014): “Delay Competition to Increase Competition: Should Reverse Payments be Banned per se?” .

Meunier, V. and J. Padilla (2015): “Should reverse payments be prohibited per se?” .

Palikot, E. F. and M. Pietola (2019): “Pay-for-Delay with Settlement Externalities,” *TSE Mimeo*.

Posner, E. A., K. E. Spier, and A. Vermeule (2010): “Divide and conquer,” *Journal of Legal Analysis*, 2, 417–471.

Segal, I. (1999): “Contracting with externalities,” *The Quarterly Journal of Economics*, 114, 337–388.

——— (2003): “Coordination and discrimination in contracting with externalities: Divide and conquer?” *Journal of Economic Theory*, 113, 147–181.

Shapiro, C. (2003): “Antitrust Limits to Patent Settlements,” *RAND Journal of Economics*, 34, 391–411.

1.7 Appendix

Proof of Proposition 1

To prove Proposition 1 we will proceed in the following steps: first, we show that in the equilibrium of the continuation game, either one entrant litigates or all entrant settle. Second, we derive optimal settlement offers and I's payoff when all entrants settle. Third, we solve the game when one entrant litigates. Fourth, the condition for the equilibrium with litigation to arise is shown. Finally, we prove the monotonicity result.

Two types of equilibria: with and without litigation

In any equilibrium of the continuation game, either one and only one entrant litigates or all entrants settle.

Due to perfectly correlated litigation outcomes, if some E_j litigates in the continuation game, the best response of E_i is to either to accept the settlement offer or to reject it and wait, because waiting saves on the litigation cost but otherwise gives the same payoff as litigation. If instead none of the rival entrants litigates, the best response of E_i is to either to accept the settlement offer or to reject it and litigate because the litigation payoff is positive by Assumption 2, whereas the payoff from waiting is zero.

We can thus cover all potential equilibria by first analyzing the ones where all entrants settle and then the ones where one entrant litigates, and the others either wait or settle.

It is convenient to denote $s(t)$ as the number of entrants who settle with an entry date t or later. At any date t, we can then categorize the entrants into three groups:

- $s(t)$ delayed entrants,
- $s(0) - s(t)$ entrants who settle and enter before t,
- $N - s(0)$ entrants who reject the settlement offer and either litigate or wait.

Equilibria without litigation

Suppose first that all entrants settle in equilibrium. Then, by definition, in the continuation game each E_i has a best response to settle given that all rival entrants settle. Thus, the payoff E_i obtains from accepting the settlement offer,

$$-p_i + \int_{t_i}^{1} \pi\left(1 + N - s\left(t\right)\right) \mathrm{d}t,$$

is no less than the payoff E_i *would* obtain by rejecting the offer and going to court instead:

$$-c+(1-\theta)\int_{1-l}^{\max\{t_i,1-l\}} \pi\left(2+N-s\left(t\right)\right)\mathrm{d}t+(1-\theta)\int_{\max\{t_i,1-l\}}^{1} \pi\left(1+N-s\left(t\right)\right)\mathrm{d}t.$$

As it is not profitable for I to leave E_i strictly better off from the settlement, the equilibrium payment $p_i\left(t_i\right)$ satisfies

$$\begin{aligned} p_i\left(t_i\right) = c + \int_{t_i}^{1} \pi\left(1+N-s\left(t\right)\right)\mathrm{d}t - (1-\theta)\int_{1-l}^{\max\{t_i,1-l\}} \pi\left(2+N-s\left(t\right)\right)\mathrm{d}t \\ - (1-\theta)\int_{\max\{t_i,1-l\}}^{1} \pi\left(1+N-s\left(t\right)\right)\mathrm{d}t. \end{aligned}$$

Note that, given the offers $\{p_i\left(t_i\right), t_i\}$ there exists a unique continuation equilibrium where all entrants accept the settlement offer. In any other equilibrium, as shown above, one entrant would litigate, whereas the others would either wait or settle. However, the litigator would have a best response to accept the settlement deal instead.

Importantly, the pay-off from the settlement relative to the pay-off from litigation weakly increases in the number of entrants who reject the settlement offer and wait instead. Thus, if it is optimal for an entrant to accept the settlement when all other entrants accept, it is also optimal to accept, when some entrants reject the settlement.

$$\frac{\mathrm{d}p_i\left(t\right)}{\mathrm{d}t} = \begin{cases} -\pi\left(1+N-s\left(t\right)\right) & \text{if } t_i < 1-l, \\ -\theta\pi\left(1+N-s\left(t\right)\right) - (1-\theta)\,\pi\left(2+N-s\left(t\right)\right) & \text{if } t_i > 1-l. \end{cases}$$

Furthermore,

$$p_i(1) = c - (1-\theta) \int_{1-l}^{1} \pi(2 + N - s(t)) \, \mathrm{d}t.$$

The total payment can then be expressed as the number of settlements times the payment associated with full entry delay, subtracting changes in payments due to early entry dates:

$$\begin{aligned}
\sum_{i=1}^{N} p_i(t_i) =& p_i(1) N - \int_0^1 (N - s(t)) \, \mathrm{d}p_i(t) \\
=& \int_0^{1-l} (N - s(t)) \pi(1 + N - s(t)) \, \mathrm{d}t \\
&+ \theta \int_{1-l}^{1} (N - s(t)) \pi(1 + N - s(t)) \, \mathrm{d}t \\
&- (1-\theta) \int_{1-l}^{1} s(t) \pi(2 + N - s(t)) \, \mathrm{d}t + Nc.
\end{aligned}$$

The payoff of I is the sum of the profits it makes in the market and the total payment. During the litigation period I has the possibility to obtain the entire industry profit:

$$\int_0^{1-l} \left[\Pi(1 + N - s(t)) + (N - s(t)) \pi(1 + N - s(t))\right] \mathrm{d}t$$

Industry profit is maximized under monopoly, thus I has no incentive to allow for entry, so $s(t) = N$ for all $t \leq 1 - l$. After the litigation period, I's profit depends on the outcome of the litigation. Adding the payments we have:

$$\begin{aligned}
&\theta \int_{1-l}^{1} \left[\Pi(1 + N - s(t)) + (N - s(t)) \pi(1 + N - s(t))\right] \mathrm{d}t \\
+ (1-\theta) &\int_{1-l}^{1} \left[\Pi(1 + N - s(t)) - s(t) \pi(2 + N - s(t))\right] \mathrm{d}t \\
&+ Nc.
\end{aligned}$$

Notice that the expected instantaneous profit depends on time only through entry dates, so the problem is linear. If the incumbent finds it profitable to delay entry until any $t > 1 - l$, it will also want to delay to $t + \epsilon > t$ and so on until patent expiry date. Consequently,

$$s(t) = \begin{cases} N & \text{for all } t \in [0, 1-l], \\ X(\theta) & \text{for all } t \in (1-l, 1], \end{cases}$$

where $X(\theta)$ maximizes

$$g(x;\theta) = \theta\left[(N-x)\pi(1+N-x) + x\pi(2+N-x)\right] \\ + \Pi(1+N-x) - x\pi(2+N-x).$$

subject to $x \leq N$. The payoff in equilibrium without litigation is then

$$S(\theta) := (1-l)\Pi(1) + lg(X(\theta);\theta) + Nc.$$

Equilibria with litigation

Suppose now that one entrant litigates, $s(0) \leq N-1$ entrants settle and $N-1-s(0)$ entrants wait in equilibrium. By offering an entrant a negative payoff from the settlement (for example a strictly positive payment with delay until patent expiration), I can be sure that the entrant rejects the deal and either litigates or waits in the continuation equilibrium. Therefore, it suffices to consider the best responses of those who settle. As one entrant litigates, by accepting the settlement offer, E_i gets

$$-p_i + \int_{t_i}^{\max\{t_i,1-l\}} \pi(1+s(0)-s(t))\,\mathrm{d}t + \theta\int_{\max\{t_i,1-l\}}^{1} \pi(1+s(0)-s(t))\,\mathrm{d}t \\ + (1-\theta)\int_{\max\{t_i,1-l\}}^{1} \pi(1+N-s(t))\,\mathrm{d}t,$$

which is no less than the payoff E_i *would* obtain by rejecting the offer and waiting instead:

$$(1-\theta)\int_{1-l}^{\max\{t_i,1-l\}} \pi(2+N-s(t))\,\mathrm{d}t + (1-\theta)\int_{\max\{t_i,1-l\}}^{1} \pi(1+N-s(t))\,\mathrm{d}t.$$

Again, as it is not profitable for I to leave E_i strictly better off from the settlement, the equilibrium payment $p_i(t_i)$ satisfies

$$
\begin{aligned}
p_i(t) = &\int_{t_i}^{\max\{t_i,1-l\}} \pi(1+s(0)-s(t))\,\mathrm{d}t + \theta \int_{\max\{t_i,1-l\}}^{1} \pi(1+s(0)-s(t))\,\mathrm{d}t \\
&- (1-\theta)\int_{1-l}^{\max\{t_i,1-l\}} \pi(2+N-s(t))\,\mathrm{d}t.
\end{aligned}
$$

Observe that, by the same logic as in the previous subsection, the offers $\{p_i(t_i), t_i\}$ imply that there exists a unique continuation equilibrium where $s(0)$ entrants accept the settlement offer. If some of them were to reject and wait instead, the payoffs from accepting the settlement deal can only increase relative to the payoffs from waiting.

Differentiating $p_i(t_i)$ with respect to t_i we obtain

$$
\frac{\mathrm{d}p_i(t_i)}{\mathrm{d}t_i} = \begin{cases} -\pi(1+s(0)-s(t)) & \text{if } t_i < 1-l \\ -\theta\pi(1+s(0)-s(t)) - (1-\theta)\pi(2+N-s(t)) & \text{if } t_i > 1-l \end{cases}
$$

and

$$
p_i(1) = -(1-\theta)\int_{1-l}^{1} \pi(2+N-s(t))\,\mathrm{d}t.
$$

As earlier, we may write:

$$
\begin{aligned}
\sum_{i=1}^{N} p_i(t_i) =& p_i(1)\,s(0) - \int_0^1 (s(0)-s(t))\,\mathrm{d}p_i(t) \\
=& \int_0^{1-l} (s(0)-s(t))\,\pi(1+s(0)-s(t))\,\mathrm{d}t \\
&+ \theta\int_{1-l}^{1} (s(0)-s(t))\,\pi(1+s(0)-s(t))\,\mathrm{d}t \\
&- (1-\theta)\int_{1-l}^{1} s(t)\,\pi(2+N-s(t))\,\mathrm{d}tc.
\end{aligned}
$$

During the litigation period I obtains the entire industry profit:

$$
\int_0^{1-l} \left[\Pi(1+s(0)-s(t)) + (s(0)-s(t))\,\pi(1+s(0)-s(t))\right]\mathrm{d}t,
$$

which is maximized under monopoly: $s(0) = s(t)$ for all $t \leq 1-l$. After the litigation period, incumbent's profit depends on the outcome of litigation:

$$\begin{aligned} &\theta \int_{1-l}^{1} \left[\Pi\left(1+s(0)-s(t)\right)+\left(s(0)-s(t)\right)\pi\left(1+s(0)-s(t)\right)\right]\mathrm{d}t \\ +&(1-\theta)\int_{1-l}^{1}\left[\Pi\left(1+N-s(t)\right)-s(t)\,\pi\left(2+N-s(t)\right)\right]\mathrm{d}t \\ &-C. \end{aligned}$$

Again, if the incumbent finds profitable to delay entry until any $t > 1-l$, she will also want to delay to $t+\epsilon > t$ and so on until patent expiry date. Consequently, the incumbent's problem reduces to selecting an entry schedule

$$s(t) = \begin{cases} Y + X & \text{for all } t \in [0, 1-l], \\ X & \text{for all } t \in (1-l, 1], \end{cases}$$

where Y and X maximize

$$\theta\left[\Pi(1+y)+y\pi(1+y)\right]+(1-\theta)\underbrace{\left[\Pi(1+N-x)-x\pi(2+N-x)\right]}_{=g(x;0)}$$

subject to $x \leq x+y \leq N-1$. As industry profit is decreasing in the number of firms in the market, $Y = 0$. Furthermore, X maximizes $g(x;0)$ and by Assumption 1 $X = X(0) \leq N-1$. Overall, the payoff in equilibrium with litigation is thus

$$L(\theta) \coloneqq (1-l+\theta l)\,\Pi(1)+(1-\theta)\,lg\left(X(0);0\right)-C.$$

Comparison

The difference between I's payoffs in equilibrium without and with litigation can be written as

$$\frac{L(\theta)-S(\theta)}{l} = \underbrace{\theta\Pi(1)+(1-\theta)\,g\left(X(0);0\right)-g\left(X(\theta);\theta\right)}_{:=f(\theta)} - \underbrace{\frac{C+Nc}{l}}_{:=\lambda}.$$

It thus follows that there exists litigation in equilibrium if $\lambda < f(\theta)$ and no litigation if the reverse inequality holds.

Monotonicity

It remains to show that $X(\cdot)$ is weakly increasing and satisfies $X(1) = N$. For the last part, note that

$$g(x;1) = \Pi(1+N-x) + (N-x)\,\pi(1+N-x)$$

is the industry profit, which is maximized at $x = N$. Furthermore, we can rewrite $g(x;\theta)$ as

$$\begin{aligned} g(x;\theta) = {} & \theta\left[\Pi(1+N-x) + (N-x)\,\pi(1+N-x)\right] \\ & + (1-\theta)\left[\Pi(1+N-x) - x\pi(2+N-x)\right], \end{aligned}$$

where the term multiplying θ is the industry profit, which is increasing in x, and the term multiplying $(1-\theta)$ is maximized at $X(0) \leq N-1 < N = X(1)$. Thus, clearly, $X(\theta) = \arg\max_x g(x;\theta)$ is weakly increasing in θ.

Proof of Proposition 2

Note that $g(x;\theta)$ defines a family of affine functions of θ parametrized by x. The epigraph of an affine function is a half-space and any intersection of half-spaces is a convex set. The value function $g(X(\theta);\theta)$ is therefore convex and piecewise linear. Notice that

$$\frac{\partial g(x;\theta)}{\partial \theta} = (N-x)\,\pi(1+N-x) + x\pi(2+N-x) > 0.$$

Therefore, $g(X(\theta);\theta)$ is a strictly increasing, convex, piecewise linear function of θ. But then

$$f(\theta) = \theta \underbrace{\Pi(1)}_{=g(X(1);1)} + (1-\theta)\,g(X(0);0) - g(X(\theta);\theta)$$

is concave, piecewise linear function of θ. Furthermore, it clearly satisfies $f(0) = f(1) = 0$, so that for any $\lambda < \max f$, we can define thresholds of patent strength, $\underline{\theta}(\lambda) < \overline{\theta}(\lambda)$, where $\underline{\theta}(\lambda)$ satisfies $\underline{\theta}(0) = 0$ and increases in λ, and $\overline{\theta}(\lambda)$ satisfies $\overline{\theta}(0) = 1$ and decreases in λ, such that there is litigation in equilibrium if and only if $\theta \in \left[\underline{\theta}(\lambda), \overline{\theta}(\lambda)\right]$.

Proof of Proposition 3

Suppose first that all entrants settle in equilibrium: $s(0) = N$. Then, for each E_i the equilibrium payment is given by

$$p_i(t_i) = c + \int_{t_i}^{1} \pi(2 + N - s(t))\,\mathrm{d}t - (1-\theta)\,l\pi(1+N),$$

where

$$\frac{\mathrm{d}p_i(t_i)}{\mathrm{d}t_i} = -\pi(1 + N - s(t))$$

and

$$p_i(1) = c - (1-\theta)\,l\pi(1+N).$$

Plugging the total payment

$$\begin{aligned}\sum_{i=1}^{N} p_i(t_i) =& p_i(1)\,N - \int_0^1 (N - s(t))\,\mathrm{d}p_i(t) \\ &= \int_0^1 (N - s(t))\,\pi(1 + N - s(t))\,\mathrm{d}t \\ &- (1-\theta)\,lN\pi(1+N) + Nc\end{aligned}$$

into the incumbent's payoff we obtain:

$$\int_0^1 \left[\Pi(1 + N - s(t)) + (N - s(t))\,\pi(1 + N - s(t))\right]\mathrm{d}t$$

less a constant term. Industry profit is maximized under monopoly, so $s(t) = N$ for all $t \leq 1$ is optimal. This gives the incumbent a payoff

$$S_c(\theta) \coloneqq \Pi(1) - N \cdot \left[(1-\theta)(1-l)\,\pi(1+N) - c\right].$$

Suppose now that at least some E_j litigates in equilibrium. Then, for each E_i who settles we have

$$p_i(t_i) = \int_{t_i}^{\max\{t_i,1-l\}} \pi(1+s(0)-s(t))\,\mathrm{d}t + \theta \int_{\max\{t_i,1-l\}}^{1} \pi(1+s(0)-s(t))\,\mathrm{d}t.$$

Differentiating with respect to t_i we obtain

$$\frac{\mathrm{d}p_i(t_i)}{\mathrm{d}t_i} = \begin{cases} -\pi(1+s(0)-s(t)) & \text{if } t_i < 1-l \\ -\theta\pi(1+s(0)-s(t)) & \text{if } t_i > 1-l \end{cases}$$

and

$$p_i(1) = 0.$$

The total payment is

$$\begin{aligned} \sum_{i=1}^{N} p_i(t_i) =& p_i(1)\,s(0) - \int_0^1 (s(0)-s(t))\,\mathrm{d}p_i(t) \\ =& \int_0^{1-l} (s(0)-s(t))\,\pi(1+s(0)-s(t))\,\mathrm{d}t \\ &+\theta \int_{1-l}^{1} (s(0)-s(t))\,\pi(1+s(0)-s(t))\,\mathrm{d}t. \end{aligned}$$

During the litigation period the incumbent has a possibility to obtain the entire industry profit:

$$\int_0^{1-l} \left[\Pi(1+s(0)-s(t)) + (s(0)-s(t))\,\pi(1+s(0)-s(t))\right]\mathrm{d}t.$$

Industry profit is maximized under a monopoly, thus the incumbent has no incentive to allow for entry, so $s(t) = s(0)$ for all $t \leq 1-l$. After the litigation period, the incumbent's profit depends on the outcome of litigation:

$$\begin{aligned} &\theta \int_{1-l}^{1} \left[\Pi(1+s(0)-s(t)) + (s(0)-s(t))\,\pi(1+s(0)-s(t))\right]\mathrm{d}t \\ &+(1-\theta)l\Pi(1+N) \\ &-C, \end{aligned}$$

which is maximized by setting $s(t) = s(0)$ for all $t > 1 - l$. Thus, in equilibrium with litigation, I obtains

$$L_c(\theta) := (1-l)\,\Pi(1) + l\,[\theta\Pi(1) + (1-\theta)\,\Pi(1+N)] - C,$$

which is always less than

$$\begin{aligned} S_c(\theta) &= \Pi(1) - N \cdot [(1-\theta)\,l\pi(1+N) - c] \\ &= (1-l)\,\Pi(1) + l\,[\theta\Pi(1) + (1-\theta)\,(\Pi(1) - N\pi(1+N))] + Nc \\ &\geq (1-l)\,\Pi(1) + l\,[\theta\Pi(1) + (1-\theta)\,\Pi(1+N)] - C = L_c(\theta)\,, \end{aligned}$$

by the assumption that the industry profit is decreasing in the number of firms in the market.

Proof of Proposition 4

Suppose first that all entrant settle in equilibrium, with $t_i \in [0, 1-l]$ for each E_i. Following similar steps as in the proof of Proposition 1, we may calculate the the payoff of I and show that during the litigation period it is never profitable for I to accommodate entry to the market. Thus, $t_i = 1 - l$ for each E_i, which yields I the payoff

$$S_0(\theta) := (1-l)\,\Pi(1) + lg(0;\theta) + Nc,$$

where

$$g(0;\theta) = \Pi(1+N) + \theta N\pi(1+N)\,.$$

Suppose now that one entrant litigates, $s(0) \leq N-1$ entrants settle and $N - 1 - s(0)$ entrants wait in equilibrium. Again, following similar steps as in the proof of Proposition 1, we may show that it is never profitable for I to license the patent. Thus, $s(0) = 0$ is optimal and I's payoff from litigation is

$$L_0(\theta) := (1 - l + \theta l)\,\Pi(1) + (1-\theta)\,lg(0;0) - C,$$

with $g(0;0) = \Pi(1+N)$. Taking the difference between I's payoffs in equilibrium without and with litigation one gets

$$\frac{L_0(\theta) - S_0(\theta)}{l} = \underbrace{\theta\Pi(1) + (1-\theta)\, g(0;0) - g(0;\theta)}_{:=f_0(\theta)} - \underbrace{\frac{C+Nc}{l}}_{:=\lambda},$$

where

$$f_0(\theta) = \theta\left[\Pi(1) - \Pi(1+N) - N\pi(1+N)\right].$$

It thus follows that there exists litigation in equilibrium if $\lambda < f(\theta)$ and no litigation if the reverse inequality holds.

Numerical example

The instantaneous profit functions are determined by a textbook Cournot quantity-setting game with zero marginal costs and an inverse demand $1-Q$, where Q denotes industry profit; firm-level outputs are $q(1) = \frac{1}{2}$, $q(2) = \frac{1}{3}$ and $q(3) = \frac{1}{4}$, resulting in equilibrium profits $\Pi(1) = \frac{1}{4}$, $\Pi(2) = \pi(2) = \frac{1}{9}$ and $\Pi(3) = \pi(3) = \frac{1}{16}$. We thus have

$$\begin{aligned}
g(0;\theta) &= \theta 2\pi(3) + \Pi(3) = \frac{2\theta+1}{16}, \\
g(1;\theta) &= \theta\left[\pi(2) + \pi(3)\right] + \Pi(2) - \pi(3) = \frac{25\theta+7}{144}, \\
g(2;\theta) &= \theta 2\pi(2) + \Pi(1) - 2\pi(2) = \frac{8\theta+1}{36},
\end{aligned}$$

so that

$$X(\theta) = \begin{cases} 0 & \text{if } \theta \leq \frac{2}{7}, \\ 1 & \text{if } \theta \in \left[\frac{2}{7}, \frac{3}{7}\right], \\ 2 & \text{if } \theta \geq \frac{3}{7}, \end{cases}$$

and

$$f(\theta) = \begin{cases} \frac{\theta}{16} & \text{if } \theta \leq \frac{2}{7}, \\ \frac{\theta+1}{72} & \text{if } \theta \in \left[\frac{2}{7}, \frac{3}{7}\right], \\ \frac{5(1-\theta)}{144} & \text{if } \theta \geq \frac{3}{7}. \end{cases}$$

In particular,

$$f\left(\frac{2}{7}\right) = \frac{1}{56},$$
$$f\left(\frac{3}{7}\right) = \frac{5}{252}.$$

Thus, applying Proposition 1, if $\lambda > f(\theta)$, both entrants obtain a license if $\theta \leq \frac{2}{7}$, one entrant obtains a license while the other one is delayed if $\theta \in \left[\frac{2}{7}, \frac{3}{7}\right]$ and both entrants are delayed if $\theta \geq \frac{3}{7}$. If instead $\lambda < f(\theta)$, one of the entrants litigates while the other one waits.

The instantaneous consumer surpluses can be calculated by integrating the demand function from zero to the equilibrium industry output:

$$CS(1) = \frac{1}{2}[q(1)]^2 = \frac{1}{8},$$
$$CS(2) = \frac{1}{2}[2q(2)]^2 = \frac{2}{9},$$
$$CS(3) = \frac{1}{2}[3q(3)]^2 = \frac{9}{32}.$$

The expected consumer surplus under litigation is

$$\theta CS(1) + (1-\theta) CS(3) = \theta \frac{1}{8} + (1-\theta)\frac{9}{32},$$

which equals $CS(2)$ at $\theta = \frac{17}{45}$. When θ is below this threshold the consumer surplus in (license, delay) is higher than in (litigate, wait).

Chapter 2

The impact of online reputation on ethnic discrimination

This chapter is based on a research project with Xavier Lambin (Lambin and Palikot (2020)).

2.1 Introduction

The online economy promised to eliminate offline frictions and facilitate collaboration among strangers. Reputation systems (reviews and ratings) provide a key mechanism for this: by aggregating information about past transactions, they discipline buyer and seller behavior and favor high quality types (e.g., Tadelis (2016)).[1] This should ensure the efficient functioning of online markets. Yet, there exists substantial evidence of severe discrimination online. On Airbnb, black hosts charge less than non-black hosts for equivalent rentals, and booking requests from black guests are less likely to be accepted (Edelman et al. (2017); Edelman and Luca (2014)). The goal of this paper is to investigate this apparent contradiction.

We collect data on a ridesharing platform that reconcile these seemingly incompatible facts. We find evidence of ethnic discrimination against minority drivers but also observe that reputation-building, thanks to passenger reviews, allows drivers to overcome *initial* discrimination. Estimating a model of career concerns, we show that the reputation system does indeed enable minority drivers to mitigate the handicap from which they initially suffer. However, building a reputation comes at a cost; as a result, the foregone payoffs stemming from the initial prejudice appear to be quantitatively important.

To perform this study, we have collected data on BlaBlaCar, a prominent French carpooling platform. BlaBlaCar is mostly used for inter-city trips with an average length of 400 km. Hence, the rides typically lead to several-hour-long interactions. Two features of the platform design are critical to our analysis. First, passengers can indeed discriminate. When searching for a ride, passengers see the profiles of all available drivers, which include their names, photos, and all the reviews from previous rides. Second, drivers set prices and collect reputation. Thus, they can act to influence demand. By exerting effort

[1] A key feature underlying the success of the “sharing economy" is the efficacy of reputation systems in building trust across social divides. See a talk by Joe Gebbia, a co-founder of Airbnb: https://www.youtube.com/watch? v=16cM-RFid9U, last accessed October 22, 2019. Furthermore, Frederic Mazzella, BlaBlaCar CEO, claims that the company’s reputation system creates “a sense of trust almost comparable to the level of trust in friends” (Mazzella and Sundararajan (2016)).

to obtain positive reviews, and setting low prices to collect these reviews at a faster pace, they can boost their reputation.

Our data show sizable differences across ethnic groups in terms of listing popularity (measured by the number of clicks they generate), the number of seats sold, and revenue generated by the listing. This disparity is robust to a rich set of driver-and listing-specific controls. Second, the gap is concentrated in the beginning of drivers' careers and shrinks as they receive reviews. Ethnic minority drivers with fewer than five reviews earn twelve percent less revenue than do nonminority entrants. This difference declines to seven percent for drivers with more than five and fewer than fifteen reviews and is statistically insignificant for users with more than forty reviews. Third, we show that the change of sample composition due to the exit of underperforming minority drivers is not the mechanism behind our results. Fourth, the analysis of the within driver variations in prices and grades reveals that drivers set lower prices and receive higher grades when they are new on the platform. Both effects are stronger for minority than nonminority drivers.

To highlight the causal link between new reviews and improvements in the economic performance of minority drivers, we exploit a natural experiment consisting of demand shocks. We carry out a difference-in-differences analysis where the treated group used the platform during an event of extraordinarily high demand caused by a railway strike, while the control group used the platform on a regular (non-strike) day. The treatment is an exogenous increase in the number of reviews available on profiles of drivers that happened to be driving on a strike day. We find that the minority drivers in the treated group achieved substantially higher revenue after the treatment than did the minority drivers in the control group.

Minority drivers have a strong incentive to build a reputation. To study how they respond to this incentive by *investing* in reputation, and to evaluate the costs of the initial prejudice, we propose a model of career concerns. Our model builds on Holmström (1999); drivers, characterized by intrinsic types (initially incompletely known) and marginal costs, set prices and exert efforts to maximize life long consumption. Passengers observe a set of available drivers and choose the one that maximizes their expected utility. They have prior beliefs about the distribution of drivers' types, which are population-specific and might be incorrect. After a ride, the passenger reports the quality of service; the report is used in successive periods to form posterior beliefs about the driver's type. The quality of service is a function of the driver's type, the amount of effort she puts in, and a random shock. Passengers observe and report the overall quality, not the individual components.

Drivers' pricing and effort decisions exhibit static-dynamic tradeoffs: they can decide to offer discounts and exert costly effort to build a high reputation quickly. The incentive to invest in reputation is strong when passengers value reputation highly, and a grade has a substantial impact on posterior beliefs; the more randomness the reviews exhibit, the lower are the efforts. Furthermore, there are decreasing returns to investing in reputation because each subsequent grade has a smaller impact on posterior beliefs. As a result, both efforts and discounts tend to zero over time.

In a market defined as a day and route combination, we observe all available drivers, their characteristics, prices, and the number of sold seats. We also know how many times each listing was viewed by potential passengers, which gives us a precise measure of the number of passengers looking for a ride and allows us to model passengers' choice problem. Each passenger chooses a driver that maximizes her expected utility from a set of available drivers and the outside option. We estimate the parameters of demand by maximizing a loglikelihood function. The crucial assumption allowing us to identify the parameters of the supply is that after a certain number of reviews, enough information is available on drivers' profiles so that in the subsequent periods they do not exert effort or offer discounts.[2] We identify drivers' types and their marginal costs from grades and prices observed after the reputation building stage.

We use market outcomes to back out beliefs about the quality of service.

[2]The model shows that efforts and discounts tend to zero as drivers collect reviews. The within driver variation in prices and grades exhibit patterns consistent with investing in reputation until approximately the tenth review.

We show that the market expects a minority driver with no reviews to be of quality 4.17 (i.e., 8th percentile of the distribution of grades) on a scale of 1 to 5 despite grading them after the trip 4.62 (48th percentile) on average. The disparity between the expected and given grades is the consequence of incorrect prior beliefs.

Prior beliefs influence incentives to invest in reputation. An additional review leads, on average, to an improvement in posterior beliefs about the quality of service of minority drivers. Consequently, minority drivers offer low introductory prices that increase the chance of selling a seat and being reviewed. The optimal prices that contain the component of investing in reputation are over eight percent lower than the price that would maximize current pay-off (the discount offered by nonminority entrants at the reputation building stage is four percent). The incentive to exert effort depends on future profits and the amount of uncertainty about the driver's type. Minority drivers initially have lower profits, but there is higher uncertainty about their types. Considering both effects, we find that they have higher incentives to exert effort than nonminority drivers.

Establishing a reputation is costly as minority drivers have to go through an initial period of low outcomes and additionally need to *invest* in reputation building. In a counterfactual, we assume that passengers have correct prior beliefs about the quality of service offered by minority drivers. We can quantify the cost of the incorrect priors and resulting discrimination by comparing the counterfactual profits to the baseline scenario: we show that the average pay-off of minority drivers over the first fifteen rides is nineteen percent higher in the counterfactual case.

In a second counterfactual, we study what happens when the initial disparity between minority and nonminority drivers does not fade away. In this scenario, passengers always consider minority drivers to be of a lower quality.[3] As a result, minority drivers' incentives to invest in reputation vanish, they increase introductory prices and exert much less effort. Their average pay-off

[3]The expected quality of service of individual minority drivers suggested by their average grades is always decreased by the size of the initial gap.

throughout the first fifteen rides is eight percent lower than the baseline.

Finally, we analyze the effects of the introduction of ethnicity-blind profiles, as proposed by Edelman et al. (2017). In this experiment, passengers are ex-ante uncertain whether a driver is from a minority or not. When passengers cannot establish the ethnicity of a driver based on the profile, there is no discrimination at the booking stage, which influences the prices and efforts of both minority and nonminority drivers. Minority drivers increase their prices and offer a better quality of service. Their profits increase substantially, nonminority drivers' profits are reduced.

Relation to literature: This paper relates to several strands of economic literature. First, the differences in economic outcomes across ethnic groups have been studied for a long time, see, e.g., Kuznets (1955); Alesina et al. (2016) show the extent of ethnic inequality worldwide. The negative impact of ethnic discrimination on economic outcomes is well documented: Banerjee and Munshi (2004) quantify the aggregate loss due to discriminatory investment decisions, and Hjort (2014) shows high economic costs of ethnic preferences in team production.[4] Discrimination against ethnic minorities in digital markets has been mostly studied in the context of short-term house rentals.[5] In the case of ridesharing, Farajallah et al. (2019) show that ethnic minority drivers set lower prices than nonminority drivers.[6] We contribute to this literature by documenting a gap in revenues and economic profits. Most importantly, we

[4]The economic theory of discrimination generally follows two approaches. Taste-based discrimination, formalized by Becker (1971), attributes discrimination to preference against interacting with some economic agents. While, the theory of statistical discrimination, due to Phelps (1972) and Arrow (1973), explains discrimination in terms of differences in the expected quality across groups; when an individual agent's quality is not observed, the expectation of it is formed based on the observed minority status. The distinction between statistical discrimination with correct and incorrect priors has recently been discussed by Bohren et al. (2019a). Bohren et al. (2019b) formalizes the theory of dynamic discrimination.

[5]See: Edelman and Luca (2014), Edelman et al. (2017), Laouenan and Rathelot (2017), and Kakar et al. (2018).

[6]The majority of empirical work in this domain identifies a disparity in prices between minority and nonminority sellers. However, a difference in prices is not necessarily due to discrimination; we show that part of it can be explained by seller heterogeneity in unobserved characteristics, for example, marginal costs.

develop a model of belief formation and updating, which allows us to estimate incorrect prior beliefs and understand their impact on the economic outcomes of minority drivers. We also show that these beliefs are updated with reviews.[7] Furthermore, the analysis of counterfactuals allows us to quantify the cost of incorrect beliefs.

Second, our structural model builds on the literature on dynamic moral hazard. We generalize the seminal model of Holmström (1999) by introducing incorrect beliefs, competition between drivers, and pricing as an additional strategic tool.[8] The structural estimation of a career concerns model using data from a reputation system is our contribution to this literature. The estimation results allow us to study drivers' reactions to discrimination. Coate and Loury (1993) and Glover et al. (2017) argue that discrimination can be a self-fulfilling prophecy. We show that conditioned on entering the market, minority drivers facing statistical discrimination with an erroneous prior exert effort and set low introductory prices to improve their future outcomes.

Third, Ge et al. (2016) show that the magnitude of discrimination depends on how early, in the booking process, the information on ethnicity becomes available. Thus, the extent of discrimination varies with the design of a marketplace. Edelman et al. (2017) discuss various policy proposals aimed at mitigating discrimination online; such policy interventions spur reactions by all market participants. Our structural model allows us to generate counterfactuals and evaluate the welfare effects of various market designs.

The rest of this paper is organized as follows: section 3.7 introduces some important features of BlaBlaCar and the data collection process. Section 2.3 provides reduced-form results. We document the output gap between minority and nonminority drivers, the analysis of which is followed by a study of the

[7] The importance of information to minority groups is shown in experimental settings by Bartoš et al. (2016) and Cui et al. (2019). Agrawal et al. (2016) provide evidence that information benefits employees from less developed countries. The additional benefit of acquiring information about new workers is explored by Pallais (2014). Sociological research has also studied the potential of reputation systems to offset trust judgments, see, e.g., Abrahao et al. (2017); Carol et al. (2019); Tjaden et al. (2018).

[8] Employer learning has been captured before by Chiappori et al. (1999) and Altonji and Pierret (2001).

effect of reputation building and a comparison of exit patterns. Next, we perform a difference-in-differences analysis exploiting a natural experiment. In section 2.4, we introduce a model of passenger choice and drivers' career concerns. Next, in section 2.5 we discuss identification assumptions and the estimation procedure. Section 2.6 presents the estimation results. Section 2.7 describes counterfactual experiments. Finally, we conclude the paper in section 2.8.

2.2 Empirical context and data collection

BlaBlaCar is an online marketplace for ridesharing that was established in 2006 in France and today operates in 22 countries, mostly in Europe, but also Mexico, India, and Brazil. The platform has over 80 million active users.[9] BlaBlaCar is particularly popular in France, where 1.5 million passengers use it every month. There are several essential differences between BlaBlaCar and ride-hailing services, such as Uber or Lyft, we discuss them in this section.

Participation in BlaBlaCar is restricted to nonprofessional drivers; this is ensured by imposing limits on the number of seats and listings drivers can offer.[10] Typically, drivers travel on a given route for for personal reasons and use the platform to cover some of the costs. BlaBlaCar is particularly popular on long routes between major cities. In our dataset, the average trip is 400 km long. Thus, a decision to travel with someone implies interacting for several hours.

Another key feature of BlaBlaCar is that drivers' set their prices. BlaBlaCar offers a suggestion that depends only on the distance and amounts to 0.062 EUR per km. Drivers typically deviate from the suggestion.[11] Figure 2.1 shows the distributions of prices on several popular routes. There is a significant degree of price dispersion within routes.

[9] https://techcrunch.com/2019/09/24/blablacar-to-acquire-online-bus-ticketing-platform-busfor/

[10] In 2019, after our sampling period, BlaBlaCar introduced BlaBlaBus, a professional bus service.

[11] The price is capped at 0.082 EUR per km, but this cap is very rarely binding.

Figure 2.1: Price dispersion on BlaBlaCar

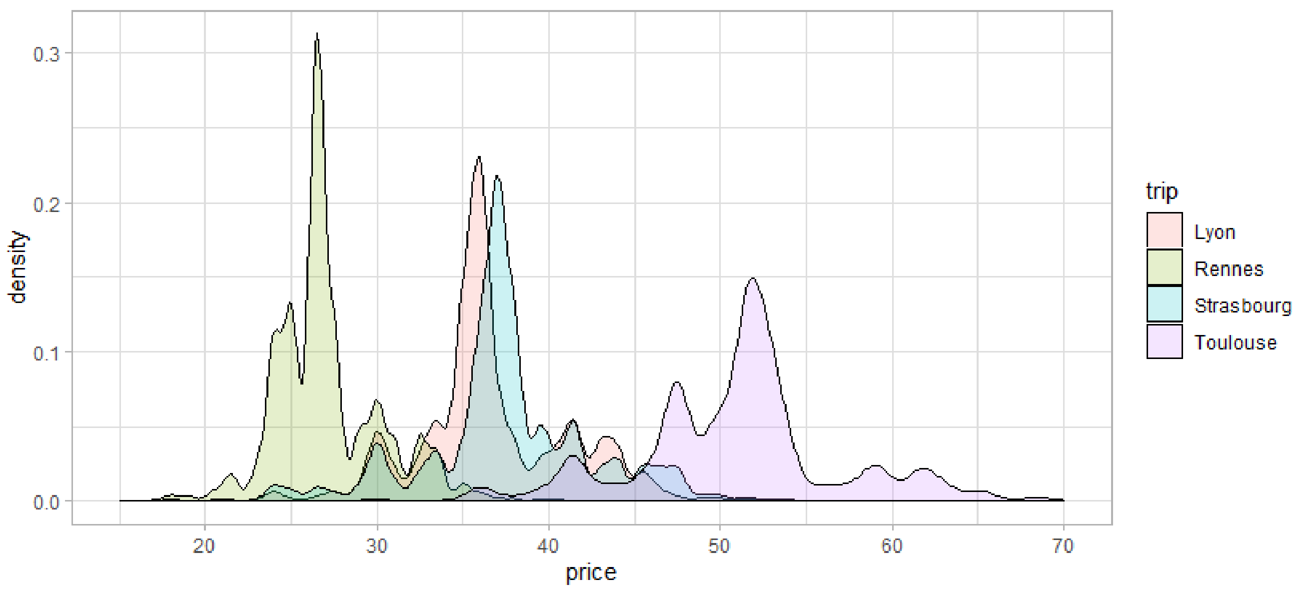

Note: Distribution of prices in euros on routes Paris to/from Lyon, Rennes, Strasbourg, Toulouse.

Before booking a ride, a potential passenger sees a list of all drivers available on a given route. By default, drivers are ranked by departure time. Some basic information is displayed at this stage: the driver's photo, name, average rating, a few details about the ride, and the price. To obtain more information, and in particular, to see the history of reviews, a prospective passenger needs to click on and visit the profile of the driver.[12] The passenger chooses the listing that she finds the most attractive and sends a booking request. Approximately half of the drivers choose the automatic acceptance feature while posting a ride; others reserve the option to reject requests. Finally, payment is made upfront via the BlaBlaCar online system. BlaBlaCar fees (see Appendix 2.9) are deducted from the price paid by the passenger.

BlaBlaCar sends multiple reminders to encourage the passengers and the driver to leave reviews. A review consists of a textual comment and a grade from 1 to 5. We have collected both the written comments and grades. We carried out a sentiment analysis of the written comments; this exercise reveals that there is a high correlation between the sentiment expressed with a written review and the associated grade. We document this in Appendix 2.9. Given this high correlation, we decided to focus only on grades. From now on, we will use *review*, *rating* and *grade* interchangeably while referring to a grade on the scale of 1 to 5.[13]

Reviews on sharing economy platforms are frequently skewed to the right (disproportionately positive). If a vast majority of reviews assign the highest possible grade, the reputation system loses its informativeness (Zervas et al. (2015) studies the implications of this). On the BlaBlaCar platform, we also see that the highest possible grade of 5 is the most popular. However, there are still enough reviews with lower grades to make the grading system meaningful. The mean grade per driver in our dataset is 4.6.

[12]Examples of profiles and listing pages are provided in Appendix 2.9.

[13]The review system has a simultaneous reveal feature, which means that a user cannot observe a received review unless she has also posted one herself or the time to write one (two weeks) has elapsed. Only after both reviews have been sent do they become available to other users. Over the years, BlaBlaCar has introduced a few changes to the reputation system, which affected grading behavior. Appendix 2.9 discusses these changes.

Data collection: We have collected our dataset using a web crawler on the website www.blablacar.fr, from 1.07.2017 to 18.03.2019. The program randomly selects a pair of cities from a predefined list of the largest cities in France and searches for available drivers. Trips start or end in Paris or its vicinity and have their other endpoints in one of the other 110 largest cities in France.

The program gathered all information accessible to prospective passengers. To do that, we open profiles of each driver available on a given route and collect all characteristics displayed on the profile, which include name, age, photo, a short biography, and the number of Facebook friends. Furthermore, we extract the entire history of received ratings and textual comments. We also observe the number of clicks and the number of sold seats for each listing. Clicking on the listing is necessary to book a trip, and a click opens a detailed description of the ride, but the passenger can still change her mind at no cost. We determine revenue by listing by calculating the product of the number of sold seats and price.

The listings that we observe have been featured on the platform for various periods of time. Some of them could have been posted just before our visit, while others could have been available for days. To account for this fact, we will control for how long a given listing is available and how many hours are left until departure.[14]

Additionally, we have matched our data with several other datasets. We establish gender and ethnicity using two complementary methods. First, we use the ethnic origins of names listed database published by the French government and supplemented with some other publicly available sources.[15] Second, we use a facial recognition software to improve our classification.[16] A detailed

[14] This explains why many of our observations have zero sold seats and zero revenue. To check whether this biases our results, for a subset of our data, we have used the BlaBlaCar API to collect the final number of sold seats and revenue. We find similar results using this additional dataset.

[15] Translations of names with foreign origins into French exhibit considerable diversity. We phonetically encode our name lists and allow for minor spelling mistakes to improve our classification.

[16] www.kairos.com

description of gender and ethnic identification is provided in Appendix 2.9, where we show that both techniques - name and facial recognition - complement each other. Our definition of minority drivers is based on names with an Arabic or African origin or connotation; in doing so, we follow most of the existing literature. However, by considering both groups and using photo recognition together with name connotation, our approach improves the practice of assigning ethnicity compared to the prior studies in this context.

We proxy the quality of the car by approximating its value by the average price of the same type of car posted on eBay in Germany. The fuel efficiency of cars is calculated by matching car models with a dataset of long-distance fuel consumption of cars. We also collect data on city-level daily average fuel prices and highway tolls to construct instrumental variables for prices. Distances and expected travel time by car or public transportation are calculated for the moment of departure using Google Maps.

We also include information specific to destination and departure cities, such as population, median income, index of crime, and a share of foreign-born residents. Additionally, we have data on strikes related to transportation services (in particular, railways) that occurred in the spring of 2018. Descriptive statistics of selected variables are shown in Table 2.1. Appendix 2.9 lists the definitions of variables and sources of supplementary data.

One hundred eight thousand drivers appear in our dataset more than once. We use these observations to construct a panel. In the panel, the median number of observations per driver is five, and the mean is 12.

We have several measures of economic performance. First, the number of clicks is our proxy for the popularity of a listing. Passengers click on many drivers before deciding with which driver to travel. The mean number of drivers that a passenger can choose from is 30. The average number of clicks that a listing received is 17. The number of clicks is also useful for capturing the number of passengers searching for a ride in a given market.

Second, we observe the number of seats sold. On average, at the point of data collection, drivers managed to sell 0.3 seats. Drivers can change the price before the first passenger books a ride, but once one seat has been sold,

Table 2.1: Descriptive statistics.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Price (EUR)	552,518	31.43	15.98	6.00	18.00	41.50	78.50
Number of clicks	536,904	16.63	17.57	0.00	3.00	25.00	77.00
Sold seats	566,023	0.26	0.58	0.00	0.00	0.00	4.00
Revenue (EUR)	559,931	6.42	15.21	0.00	0.00	0.00	82.50
Minority	566,023	0.14	0.35	0.00	0.00	0.00	1.00
Male	552,530	0.73	0.44	0.00	0.00	1.00	1.00
Driver age	558,032	37.51	12.80	18.00	27.00	47.00	68.00
Number of reviews	560,331	37.12	60.71	0.00	4.00	42.00	421.00
Published rides (total)	537,681	38.84	49.29	0.00	7.00	50.00	256.00
Reputation	516,021	4.60	0.31	1.00	4.50	4.80	5.00
Seniority (months)	559,890	44.66	28.03	1.00	23.00	64.00	118.00
Posts per month	555,962	1.44	2.17	0.01	0.26	1.62	17.24
Photo	566,023	0.97	0.18	0.00	1.00	1.00	1.00
Bio (# words)	537,475	7.44	10.38	0.00	2.00	12.00	42.00
Car value (thousands of EUR)	471,117	6.08	5.04	0.60	3.10	8.06	24.40
Fuel consumption	486,604	5.00	0.77	3.65	4.39	5.39	7.50
Automatic acceptance	566,023	0.42	0.49	0.00	0.00	1.00	1.00
Hours until departure	508,754	95.50	107.47	0.001	20.96	126.47	501.69
Posted since	560,361	5.88	7.50	0.00	1.53	6.82	52.56
Travel time by public transport	545,200	3.97	2.42	0.14	2.25	5.41	15.24
Trip length (km)	550,118	396.34	192.27	67.32	232.00	491.68	906.46
Travel cost (fuel & tolls, EUR)	458,018	57.01	29.10	0.00	33.71	72.13	142.14
Train strike	566,023	0.04	0.19	0.00	0.00	0.00	1.00
Ride description (number of words)	509,243	13.49	14.60	2.00	2.00	22.00	93.00
Median revenue (city)	532,526	18.98	2.13	13.06	17.76	20.20	30.90
weekday	566,024	0.67	0.47	0	0	1	1
luggage size	116,982	0.89	0.31	0.00	1.00	1.00	1.00
detour	116,454	0.75	0.43	0.00	0.00	1.00	1.00
allows pets	223,774	0.22	0.41	0.00	0.00	0.00	1.00

Note: See Appendix 2.9 for the definitions of variables and sources of supplementary data.

the price remains the same. Hence, all passengers pay the same price. Third, the product of a price and the number of sold seats is revenue. In the structural model, we recover marginal costs; thus, we will also be able to measure economic profits.[17]

2.3 Reduced-form evidence

This section establishes several facts about the economic outcomes of minority drivers and the impact of the reputation. First, we show the disparity in the number of clicks, sold seats, and revenue between minority and nonminor-

[17] Our dataset may miss some very successful rides that were no longer displayed when data were collected, which would lead to bias if the speed at which listings fill differs between minority and nonminority drivers. In Appendix 2.9, we explore this issue and show that its magnitude is most likely not significant. However, as the most popular listings might be those of nonminority drivers, our estimates of the output gap should be regarded as a lower bound.

ity drivers. This gap is initially quite substantial, but it decreases as drivers receive reviews. The reduction of the disparity is not due to the exit of underperforming minority drivers but to the causal impact of reviews. Second, we show trajectories of grades and prices across drivers' careers to argue that they act strategically to improve their performance. Most drivers offer low introductory prices and receive high grades when they enter the platform, and both effects are more pronounced for minority drivers. Finally, we present a natural experiment to stress the causal impact of reviews on improvement in economic outcomes of minority drivers.

2.3.1 Ethnic discrimination and the impact of reputation

A quick examination of the dataset reveals that minority drivers achieve lower outcomes than do nonminority drivers. The raw data show that despite setting on average lower prices per passenger (30.1 EUR vs. 31.6 EUR), minority drivers receive fewer clicks (15.4 vs. 16.8), sell fewer seats (0.258 vs. 0.263), and as a result earn lower revenue (5.81 EUR vs. 6.53 EUR).

Market-specific effects and other observed characteristics of drivers could explain these differences. We will now control for all variables available in our dataset. Throughout the paper, subscript i refers to drivers. We estimate the following model:

$$y_{itr} = \alpha + X_{it}\beta + Z_i\gamma + \tau_t + \xi_r + \epsilon_{itr}, \tag{2.1}$$

where t represents time, r corresponds to a route; y_{itr} is the variable of interest (i.e., the number of clicks or sold seats or the revenue), α is an intercept, X_{it} is a vector of time-varying explanatory variables , Z_i are time-invariant explanatory variables, τ_t denotes time effects, ξ_r is an effect specific to a route (a pair of cities), and ϵ_{itr} is the error term.

Table 2.2 presents the estimation results. The dependent variable in the first regression is the number of clicks; it is the number of sold seats in the second regression and revenue in the last one.

Table 2.2: Output measures regressed on driver and ride characteristics.

	Dependent variable:		
	Number of clicks	Sold seats	Revenue
Minority	−0.444*** (0.082)	−0.017*** (0.003)	−0.588*** (0.079)
Number of reviews	0.033*** (0.001)	0.002*** (0.0001)	0.041*** (0.001)
(Number of reviews)^2	−0.0001*** (<0.0001)	−0.00000*** (<0.0001)	−0.0001*** (<0.0001)
Male	−1.400*** (0.064)	0.002 (0.002)	−0.094 (0.061)
Driver age	−0.058*** (0.002)	−0.001*** (0.0001)	−0.022*** (0.002)
Posts per month	−0.557*** (0.020)	−0.010*** (0.001)	−0.201*** (0.019)
Bio (number of words)	0.001 (0.003)	0.0001 (0.0001)	0.006** (0.003)
Car value	0.006 (0.006)	−0.0001 (0.0002)	−0.010* (0.005)
Seniority (number of months months)	−0.017*** (0.001)	−0.0004*** (<0.0001)	−0.010*** (0.001)
Photo	0.799*** (0.170)	0.001 (0.006)	0.061 (0.163)
Automatic acceptance	−0.773*** (0.060)	0.131*** (0.002)	3.135*** (0.057)
Hours until departure	−0.039*** (0.0003)	−0.001*** (0.00001)	−0.021*** (0.0003)
Posted since	1.269*** (0.005)	0.011*** (0.0002)	0.292*** (0.004)
Travel time by public transport	1.080*** (0.314)	0.018 (0.011)	−1.519*** (0.299)
Length (# km)	0.007*** (0.001)	−0.0002*** (0.0001)	0.013*** (0.001)
Train strike	4.795*** (0.201)	0.128*** (0.007)	2.949*** (0.191)
Ride description (number of words)	0.033*** (0.002)	0.001*** (0.0001)	0.021*** (0.002)
Constant	13.299*** (0.588)	0.321*** (0.021)	5.671*** (0.560)
Time fixed effects	X	X	X
Route fixed effects	X	X	X
Observations	302,645	317,643	314,361
R^2	0.247	0.075	0.075
Note:			*p<0.1; **p<0.05; ***p<0.01

First, minority status has a negative coefficient and is highly statistically significant for all measures of economic performance. Second, the number of reviews has a positive impact and is highly statistically significant in all regressions. Note that increasing the number of reviews benefits both minority and nonminority drivers. The negative coefficients associated with the squared number of reviews suggest decreasing returns to accumulating reviews. Finally, younger drivers with rides that include extended descriptions experience better economic outcomes. After we control for the number of reviews, seniority on the platform has a negative coefficient.[18]

Reputation effect: When a driver has no reviews, passengers have to rely entirely on socioeconomic characteristics (age, gender, and ethnicity) to form beliefs about the expected quality of service. As the driver uses the platform, reviews left by past passengers become available on her profile and reveal individual information about the driver. As a consequence, the role of socioeconomic characteristics diminishes as the driver collects reviews.

[18] In Appendix 2.9, we control for price in a regression that uses the number of sold seats as the dependent variable; we also instrument prices with cost shifters to address the endogeneity of price and quantity.

Figure 2.2: Gap between minority and nonminority drivers decreases with reviews

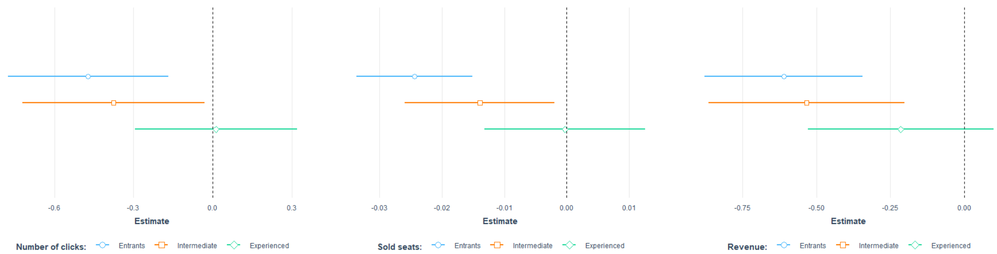

Note: Impact of minority status on: number of clicks (left), sold seats (center), revenue (right) across reputation levels: blue-entrants, orange-intermediate, green - experienced. Coefficients from OLS regressions.

If initial discrimination is due to incorrect beliefs about the expected quality of service provided by minority drivers, the intergroup disparity in economic performance will decline as individual information becomes available. This is so because reviews reveal, on average higher quality than expected ex-ante.[19]

To study the impact of reputation, we divide drivers in our dataset into three categories: entrants, defined as drivers with five or fewer reviews, intermediate (with between 6 and 15 reviews), and experienced (with more than 40 reviews). We are interested in measuring the disparity between minority and nonminority drivers in each of these groups. We estimate standard OLS regressions with the same set of controls as in Table 2.2 for drivers with different levels of experience. The full results are presented in Appendix 2.9; here, we focus on the impact of minority status only.

Figure 2.2 shows the impact of the minority status on the number of clicks (the left panel), the number of sold seats (the center panel), and revenue (the right panel) across various levels of reputation. For entrants (blue), minority status is associated with fewer clicks, sold seats, and lower revenue. The disparity between minority and nonminority drivers decreases with accumulating reviews; it is already smaller at the intermediate level of reputation, and there

[19]In contrast, if discrimination is taste-based, the information about the quality of service provided by minority drivers will not matter. In the taste-based discrimination case, the only relevant information is the ethnicity status itself.

is no statistically significant difference for drivers with more than 40 reviews.

Controlling for other observables, the initial gap in revenue (for drivers with 0 to 5 reviews) is 11.8%. It decreases to 6.9% for intermediate drivers (with 6 to 15 reviews) and is as low as 1.6% for experienced drivers (with more than 40 reviews). The results are similar for other measures of performance.[20]

Is the reputation effect due to selection? The evolution of the population of drivers on BlaBlaCar is characterized by frequent entries and exits. However, minority entrants are not more likely to quit than are nonminority entrants. The share of minority drivers is 14.6% among entrants, 13.2% in the intermediate group, and 15.6% in the experienced group. The share is relatively stable or even increasing, which suggests that selection cannot explain the reputation effect.

To provide further evidence that selection is not the mechanism behind the reduction of the disparity, in December 2018, we revisited profiles of drivers that appeared in our dataset earlier and collected their newly received reviews. The new data allow us to analyze usage intensity. We define two variables to measure the inactivity of drivers. Variable *exit* takes the value one if no new reviews were received between the last time a given driver appeared in the dataset and December 2018 and is zero otherwise. We also introduce a variable called *disaffection*, which takes the value one if the driver gathered fewer than five new reviews. Table 2.3 shows the results of the estimation of a logit model.

First, minority drivers are more likely to continue using the platform. Second, new drivers are, generally, more likely to quit. However, we find no evidence that minority entrants are leaving the platform more frequently than nonminority entrants.[21]

These results suggest that the *reputation effect* is due not to a change

[20]The gap in the number of clicks is 2.8% for entrant drivers, 2.2% for intermediate drivers, and 0.1% for experienced ones. As to the number of sold seats, the initial gap is 12.2%. It declines to 5.5% with five to fifteen reviews and to 0.1% for experienced drivers.

[21]The same analysis using the number of listings published (instead of the number of reviews collected) as a proxy for activity on the platform gives similar results.

Table 2.3: Minority entrants are not more likely to exit the platform

	Dependent variable:	
	exit	disaffection
Minority	-0.129^{***} (0.028)	-0.097^{***} (0.030)
Entrant	1.350^{***} (0.024)	1.419^{***} (0.025)
Minority*Entrant	0.079 (0.065)	0.065 (0.066)
Age	-0.005^{***} (0.001)	-0.003^{***} (0.001)
Male	-0.098^{***} (0.018)	-0.084^{***} (0.019)
Seniority (number of months)	-0.005^{***} (0.0003)	-0.005^{***} (0.0004)
Posts per month	-0.731^{***} (0.010)	-0.736^{***} (0.011)
Bio (number of words)	-0.007^{***} (0.001)	-0.007^{***} (0.001)
Constant	-0.867^{***} (0.053)	-1.377^{***} (0.058)
Other driver characteristics	X	X
Time fixed effects	X	X
Observations	160,923	160,923

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

Note: Logit regressions, exit and disaffection as dependent variables.

in the composition of the sample, but to a causal impact of reviews. We do not observe, or model opportunity costs guiding drivers' entry and exit decisions. However, our findings are consistent with the idea that drivers are aware of the *reputation effect*. They realize that after a couple of periods of underachivement, their outcomes will improve; thus, they do not leave the platform despite facing initial discrimination: although the frequent exit of entrants is an essential aspect of the dynamics of the population of drivers on BlaBlaCar, the distinction between exit rates of minority and nonminority entrants is inconsequential.

2.3.2 Strategic behavior of drivers

Establishing reputation benefits all drivers, but is particularly valuable for minority drivers. In this section, we document how drivers respond to the incentive of acquiring a reputation. We explore two dimensions - efforts put into receiving higher grades and prices chosen by the drivers.

Arguably initial reviews are more consequential as they shift the posterior belief about quality to a larger extent. Therefore, if reviews reflect efforts exerted by drivers, the initial grades should be higher than the later ones - Figure 2.3 shows that this is the case. We restrict our attention to drivers who stayed on the platform at least until they obtained 30 reviews, and we explore the variation within their grades. Thus, survivorship bias does not influence the results. Figure 2.3 shows that drivers obtain, on average higher grades when they are new; the average grade decreases until the 10th review, at which point it stabilizes.

Figure 2.3: Initial grades are higher

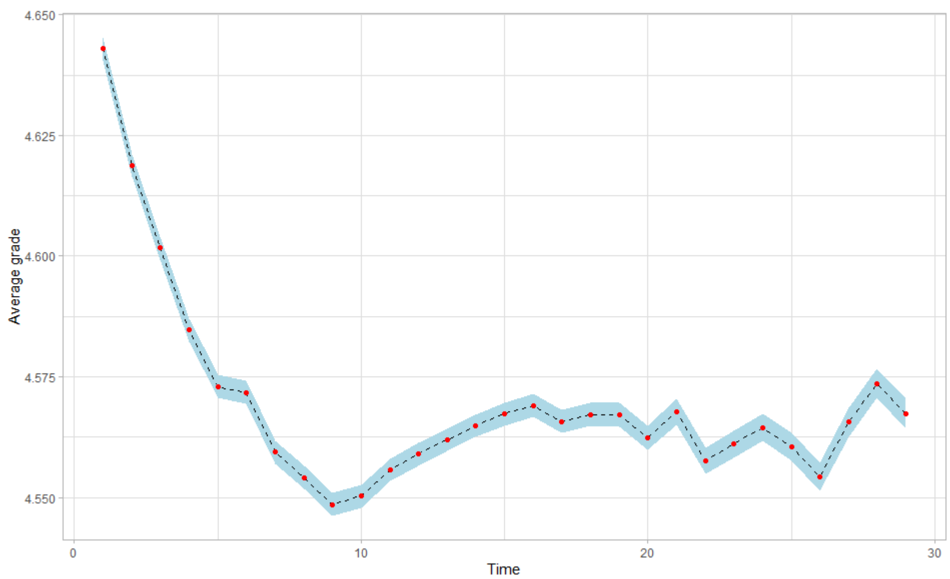

Note: The average grade from the first to 30th. Subset of drivers who used the platfrom at least until obtaining 30 reviews.

The extent to which the initial grades are higher varies across ethnic groups. In Figure 2.4, we show the difference between the early grades and the average grade a driver received after the 15th review. The gradual decrease in grades of minority drivers is more substantial than that for nonminority drivers. We interpret this as evidence of a larger effort that minority drivers exert to build a reputation.

Another way to boost reputation is to offer low introductory prices, which increases the chances of selling a seat and being reviewed. Minority drivers have an additional gain from accumulating reviews because they are, on average of higher quality than what the market expects. In Table 2.4 we present

Figure 2.4: Minority drivers exert higher efforts

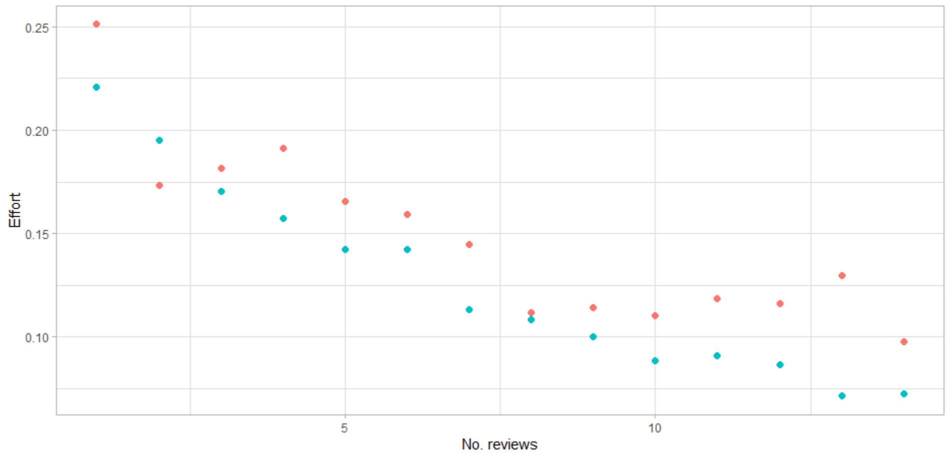

Note: Average early grades standardized by average late grades. Red dots: minority drivers, Blue dots: nonminority.

results of the estimation of within driver price variation.

We find that all drivers offer low introductory prices. The first few reviews lead to a significant increase in prices: the third review leads to an increase of 50 cents on average and the 5th review to an increase of 70 cents. However, there are decreasing returns from reviews. There is already no additional gain from the 9th review onwards. The last column of Table 2.4 introduces a distinction between minority and nonminority drivers. We observe that minority drivers set significantly lower prices when they have very few reviews; however, this effect disappears as soon as they have at least three reviews.

2.3.3 Railway strike as a quasi-experiment

The final piece of reduced-form evidence that we provide aims at highlighting a causal relationship between increasing reviews and improving economic outcomes of minority drivers.

The number of reviews that a driver has depends on the success in selling seats in the previous periods. In the data collection process, we tried to gather all information about drivers that is available to passengers, so we

Table 2.4: Within driver price variation: the impact of reputation

	Dependent variable:	
	price	
	(1)	(2)
reviews:1-2	0.307 (0.221)	0.455* (0.239)
reviews:3-4	0.495** (0.235)	0.631** (0.254)
reviews:5-8	0.691*** (0.241)	0.772*** (0.261)
reviews:9-12	0.910*** (0.260)	1.040*** (0.282)
reviews:13-16	0.798*** (0.280)	0.901*** (0.304)
reviews:17-20	0.857*** (0.303)	0.971*** (0.330)
reviews:1-2*minority		−1.040* (0.632)
reviews:3-4*minority		−0.987 (0.676)
reviews:5-8*minority		−0.610 (0.688)
reviews:9-12*minority		−0.921 (0.731)
reviews:13-16*minority		−0.747 (0.783)
reviews:17-20*minority		−0.818 (0.835)
Ride controls	x	x
Driver FE	x	x
Day FE	x	x
Observations	78,903	78,903
R^2	0.658	0.658

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

Note: Within driver variation in prices, panel estimation. Reviews are binned and used as levels.

would not need to worry about unobserved demand-relevant driver characteristics. However, some features of profiles, namely, the visual content of a photo and substance of a driver's description, are hard to capture with a proxy. If these are important to passengers, they will be correlated with the number of reviews and will bias our results.[22]

To confirm a causal relationship between the increase in the number of reviews and a reduction in the minority performance gap, we exploit a natural experiment. During our sample period, French railway workers went on a national strike.[23] The strike was organized as a sequence of two days of disruptions every five days for three months. BlaBlaCar and railways are in direct competition. Thus, a negative supply shock happening on the railway market transmits to BlaBlaCar as a positive demand shock. In April 2018, 5 million passengers traveled on BlaBlaCar, up from an average of 1.5 million.

[22]This problem is also addressed with panel estimators in Appendix 2.9.

[23]Apart from other reasons, the opposition to the plans to liberalize the European railway market and in particular to open the French market to competition was the cause of the strike.

The number of booking requests increased sixfold.[24]

All drivers, including minority drivers, faced significantly higher demand during the strike days. Figures 2.5a and 2.5b show an increase in the number of sold seats and revenue earned during the days of the strike.

Figure 2.5: Railway strike as a demand shock.

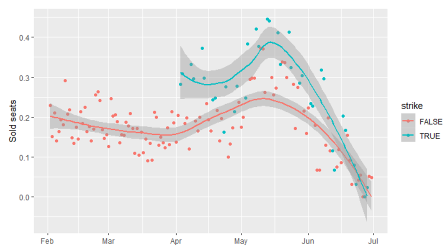

(a) Number of sold seats

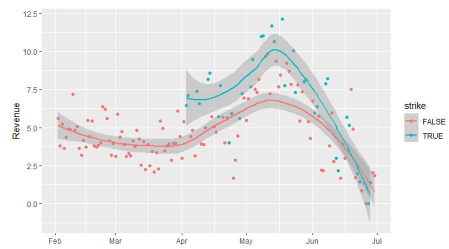

(b) Revenue

Note: Horizontal axes time; red dots days without strike; blue dots days of strike.

We will interpret the strike as a natural experiment, where the treatment is an increase in the number of reviews. This increase is due to extraordinarily high demand conditions, leading to a higher number of sold seats. The critical assumption is that drivers did not select into treatment so that the increase in the number of reviews was exogenous. We argue that BlaBlaCar drivers are not professional drivers; they travel on a given route for other reasons and do not change their plans in response to a demand shock.

To support the assumption that treatment is exogenous to driver-specific characteristics, we compare the drivers on days with and without a strike. First, selection would result in an increased number of entrants traveling on the day of the strike. Figure 2.6a shows that there is no significant difference in the number of entrants on the days of strike and non-strike days. Second, minority drivers could be aware that it is easier to sell seats on a strike-day and be more inclined to post a ride. Figure 2.6b compares the share of minority

[24] Source: www.lemonde.fr/economie/article/2018/04/03/les-transports-alternatifs-grands-gagnants-de-la-greve-a-la-sncf_5279932_3234.html

drivers on strike and non-strike days; we do not observe increased entry of minorities on strike days. During strike days, 14.7% of drivers were minority drivers. On a non-strike day, in this period, the share was 14.8%. Table 2.14 in Appendix 2.9 compares other characteristics.

Figure 2.6: No selection to treatment.

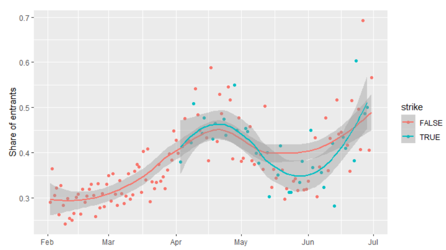

(a) Share of entrants

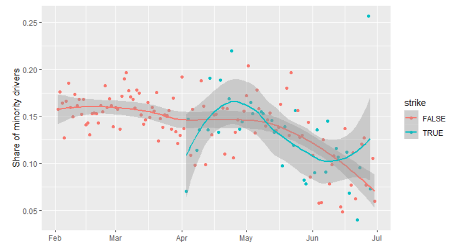

(b) Share of minority drivers

Note: Horizontal axes time; red dots days without strike; blue dots days of strike.

To show the impact of exogenous variation in the number of reviews on outcomes of ethnic minority drivers, we perform a difference-in-differences analysis. Our *treated* group represents minority drivers who happened to travel on a day of the strike;[25] $after$ indicates the period after the strikes. Finally, DiD is a product of *treated* and $after$. We estimate the following regression:

$$y_{itr} = \alpha + \beta_1 \times treated_i + \beta_2 \times after_i + \beta_3 \times DiD_i + \mathbf{X}_{it}\gamma + \mathbf{Z}_i\theta + \xi_r + \epsilon_{itr}, \tag{2.2}$$

where y_{itr} is revenue or the number of sold seats by a driver i in period t on route r, and the variable *treated* captures possible differences between the treatment and control groups prior to the demand shock. The time dummy $after$ controls for aggregate factors that would cause changes in y_{itr} even in the absence of a policy change. We are interested in the coefficient β_3 associated with the treated group in the period after the treatment. In the estimation,

[25]Some drivers used the platform multiple times during the strike, so were treated more than once.

we remove the period of strikes.

Table 2.5 presents the estimation results with revenue as a dependent variable.[26] In Table 2.15 (Appendix), we show results of the same regression, obtained if the treated group contained only minority drivers who, when driving during the strike, had fewer than three reviews. We obtain similar results with higher statistical significance.

Table 2.5: Difference-in-differences estimation with revenue as the dependent variable

	Dependent variable: revenue		
	(1)	(2)	(3)
Treated	-0.454 (0.346)	-0.419 (0.415)	-0.371 (0.414)
After	-3.603 (3.671)	-3.814 (4.019)	-4.006 (4.017)
DiD (Treated*After)	1.179* (0.645)	1.390* (0.711)	1.372* (0.711)
Minority	-0.542*** (0.086)	-0.367*** (0.098)	-0.304*** (0.098)
Driver characteristics			x
Listing characteristics		x	x
Route effects	x	x	x
Time effects	x	x	x
Observations	297,189	240,656	240,656

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

Note: Treated- minority driver who drove during the day of strike. After- period after the end of strikes.

We observe that there is no significant change in the overall revenue in the period after the strikes; the treated group appears not to differ from the control group (after the minority dummy has been included). The variable of interest is DiD - we observe that it is positive and significant across all specifications; DiD captures two effects: the correction of beliefs about the quality of treated minority drivers, and an increase in the number of reviews. Thus, its magnitude is much higher than the impact of the minority status alone. Note that a driver typically receives more than one review for each completed trip, and some drivers used the platform several times during strikes.

We believe that this result provides us with evidence of a causal link between reviews and the improvement of performance of minority drivers.

[26] The results obtained using the number of sold seats are presented in Appendix 2.9 in table 2.15.

Summary of reduced-form results: Minority drivers achieve, on average, lower economic outcomes than do nonminority drivers. The difference is statistically significant and has a substantial magnitude. This disparity is particularly pronounced for drivers with no reputation and significantly narrows down as drivers collect reviews. This effect is not due to a change in sample composition. Drivers act strategically to build reputations- they offer low introductory prices and exert efforts to get high grades; both effects are stronger for minority drivers. A difference-in-differences study using a natural experiment suggests a causal relationship between building reputation and reducing the ethnic disparity.

2.4 Model of driver's career concerns

Reputation is a valuable asset for any driver. A high average rating creates an expectation of a high quality of service. On the other hand, each additional review reduces the uncertainty about the outcome of a transaction. Therefore, minority and nonminority drivers alike have an incentive to *invest* in their reputations. However, as we argue in this paper, this incentive differs across ethnic groups. Minority drivers face bias against them, which they can correct with reviews; thus, they have a larger benefit from building reputations.

To understand how minority drivers respond to discrimination, we would ideally randomly choose some drivers to be discriminated against and compare their behavior with that of a control group that does not face discrimination but is otherwise identical. We would run this experiment long enough to understand the full dynamics of discrimination. Such an experiment is impossible. Therefore, to understand the behavior of minority drivers, we propose a model. The model will help us analyze how minority drivers set prices and vary their efforts in order to maximize their lifelong consumption.

Passengers in our model are nonstrategic players. A passenger observes available drivers and chooses the one that he or she *believes* will maximize the passenger's utility. The utility of passengers depends on the quality of service provided by the driver, the price, and other listing-specific characteristics;

quality is incompletely known, and passengers hold beliefs about it. We use the framework of statistical discrimination with a potentially incorrect prior, to understand how these beliefs are formed and updated. The belief about the quality of service provided by driver i, from population m (minority or nonminority) in period t is a function of the *prior belief* about the distribution of types in population m, and grades already available on the driver's profile.

We treat drivers as strategic players; they set prices and exert efforts to maximize their lifelong consumption. Each driver i is characterized by two unobserved characteristics: marginal cost c_i and intrinsic quality type η_i. The driver's type is initially incompletely known, and the market learns about it through reviews. A review is a truthful report of quality and depends on the driver type, the effort she exerts, and a random disturbance.

We assume following timing of the game:

1. Drivers observe characteristics of their competitors and set prices to maximize the discounted sum of future consumption.

2. Passengers looking for rides choose drivers that maximize the passengers' expected utility.

3. Drivers choose the level of effort, considering the impact of the grade obtained in period t on future consumption, and the cost of exerting effort.

4. Passengers observe a realization of quality and report on it with reviews.

2.4.1 Passengers' choice problem

A passenger j observes all available drivers and forms an expectation of utility associated with traveling with each of them. We assume that the utility is linear in characteristics of drivers, and all passengers have the same valuation for them. The expected utility of passenger j resulting from traveling with driver i from population m in period t is written as

$$\mathbf{E}\left[u_{ijtm}|w^{it}\right] = \alpha\mathbf{E}\left[w_{ijtm}|w^{it}\right] + \gamma p_{it} + \beta r_{it} + \theta\mathbf{X_{it}} + \varepsilon_{ij}, \tag{2.3}$$

where $\mathbf{E}[w_{ijtm}|w^{it}]$ stands for the expected quality given the history of reviews w^{it}, p_{it} is price, r_{it} is the number of reviews, and $\mathbf{X_{it}}$ measures other listing-specific characteristics. The passenger chooses between N available drivers or the outside option, the utility of which we normalize to zero. Passenger j chooses driver i, which we denote by $d_{ijt} = 1$, if

$$\mathbf{E}\left[u_{ijtm}|w^{it}\right] = \max\left\{\mathbf{E}\left[u_{kjtm}|w^{kt}\right], 0\right\} \ \forall k \in N.$$

Belief formation

Variable $\mathbf{E}[w_{ijtm}|w^{it}]$ summarizes what passenger j believes to be the quality of service provided by driver i in period t. We focus on two aspects of how this belief is formed. First, there is a prior belief about the distribution of types in population m, which determines the expected quality of service of drivers with no reputation and serves as a starting point for learning about the types of individual drivers. Second, the belief is updated when reviews become available.

We propose a Bayesian model of belief formation and updating to analyze a passenger's learning process. Let η_i be a measure of the driver's talent that is initially incompletely known both to the market and to the driver (the information structure is symmetric). The driver has an observable characteristic that allows the market to learn that she belongs to population m. The market and the driver have initial beliefs about the distribution of η in population m. Specifically, both the driver and the market believe that quality is distributed normally with precision (the inverse of the variance) given by h_m. The driver knows that the mean of the distribution is at $\hat{\mu}_m$. In contrast, the market believes the mean of the population to be at μ_m. The two beliefs might not coincide. Over time, the market learns about η_i by observing reviews that driver i receives. A review w_{imt} is a report of quality observed by a passenger in period t. We assume that all passengers observe and report the quality identically, so we drop the subscript j. Suppose that the quality has the following structure:

$$w_{imt} = \eta_i + a_{imt} + \epsilon_{imt}, \tag{2.4}$$

where $a_{imt} \in [0, \infty]$ is the (unobservable) effort exerted by driver i, and ϵ_{imt} is a stochastic disturbance. Passengers observe and report the total quality - w_{imt} but cannot discern its individual components.

To make an inference from observing the reviews, we need to make an assumption on the distribution of the disturbance term. Let ϵ_{imt} be distributed normally with mean zero and precision h_ϵ. It is also assumed to be independent across time.

Passengers are aware that part of the quality they observe arises from the effort put of the driver. They have a belief about the optimal level of effort a_{imt} that driver i should be exerting.[27] Therefore, observing w_{imt} will in equilibrium be equivalent to observing

$$z_{imt} \equiv \eta_i + \epsilon_{imt} = w_{imt} - a_{imt}.$$

The expectation of the type η_i of a driver from population m with the history of grades w^{it} is written as[28]:

$$\mathbf{E}\left[\eta_i | w^{it}\right] = \frac{h_m \mu_m}{h_i + t h_\epsilon} + \frac{h_\epsilon}{h_m + t h_\epsilon} \sum_{s=1}^{t} z_{ims}. \tag{2.5}$$

The expected quality in period t, from equation 2.4, is formed based on the posterior belief about the type of the driver and the expected level of effort. The expected quality is given by equation 2.6

$$\mathbf{E}\left[w_{imt} | w^{it}\right] = \frac{h_m \mu_{im}}{h_i + t h_\epsilon} + \frac{h_\epsilon}{h_i + t h_\epsilon} \sum_{s=1}^{t} (w_{ims} - a_{ims}) + a_{imt}. \tag{2.6}$$

[27] Note that, the effort is correctly anticipated, given the beliefs about the prior distribution of types. Thus, the optimal level of effort a_{imt} perceived by the market might be different from the optimal level of effort of the driver.

[28] See DeGroot (2005) for details of this derivation. Appendix 2.9 discusses the case of discrete grades.

The expected quality of driver i from population m is a weighted sum of the prior belief and the obtained grades. Three types of factors influence this expectation. The first is the population-specific prior belief about the distribution of types and the variance of types within the population. The higher the prior belief about the mean of the distribution is, the higher the posterior belief. Increasing the variance of types in the population m decreases the weight given to the prior. In other words, the mean is less informative when the population is highly dispersed. The second entails the influence of the grades available on the driver's profile relative to the prior. The higher are the grades, the higher is the posterior belief. Furthermore, when more grades are available, the lower the weight put on the prior. Finally, we can measure the informativeness of the reputation systems by the variance of the error term ϵ_{imt}. The higher the variance is, the less informative, the grades are. In the limit case of the variance equal to zero, one grade is enough to reveal the driver's type.

2.4.2 Drivers' strategic decisions

We model the behavior of drivers by assuming that they play a dynamic game of incomplete information, where their strategic choices include setting prices and exerting effort.[29] Our model is a generalization of the canonical model of career concerns of Holmström (1999). In contrast to Holmström (1999), we allow for elastic demand, include the pricing stage, introduce the populations of drivers with different distributions of types, and allow incorrect prior beliefs about these distributions.

Formally, drivers solve the problem described in equation 2.7:

$$\max_{p_{imt},a_{imt}} \mathbf{E}\left\{\sum_{t=1}^{\infty}\beta^{t-1}\left[\pi_{imt}(p_{imt},X_{imt},w^{it},c_i,\mathcal{S}_t)-g(a_{imt})\right]\right\}, \tag{2.7}$$

[29] In our approach, we do not analyze the entry decisions of drivers. This choice is motivated by the analysis of exit decisions in section 2.3. The key argument is that there is no significant difference in exit patterns between minority and nonminority entrants. We argue that minority drivers are aware that they will face bias against them; thus, we decide to focus on their response in the platform.

where π_{imt} is the instantaneous profit of driver i in period t that depends on the price p_{imt}, driver characteristics X_{imt}, the history of her grades w^{it}, the driver's marginal cost c_i and the market structure $\mathcal{S}_t$, which summarizes other drivers, namely, their characteristics and prices. The cost of effort is given by $g(a_{imt})$; we assume that it is increasing and convex. We make three additional assumptions.

Assumption 1: The level of effort is noncontractible.
Assumption 1 implies that the driver's choice of effort in period t is determined only by the impact of current efforts on consumption in periods from $t+1$ onwards.

Assumption 2: $\mathcal{S}_{t+1}$ is not a function of p_{imt} and a_{imt} for all i, t, and m.
Assumption 2 means that drivers cannot influence the market structure with their pricing and efforts. In our context, this assumption is natural because drivers rarely compete against each other more than once.

Assumption 3: Drivers consider only observed market structures (the number of potential passengers and characteristics of other drivers) as potential future market structures.
Assumption 3 implies that the drivers' expectation of future $\mathcal{S}_s$ for $s > t$ is the current market structure $\mathcal{S}_t$.

Profit-maximizing level of effort

To characterize the optimal choice of effort, we compute the derivative of equation 2.7 with respect to effort a_{imt}. We obtain the following first-order condition:[30]

$$\sum_{s=t}^{+\infty} \beta^{s-t} \times \mathbf{E}\left[\frac{\partial \pi_{ims}}{\partial a_{imt}}\right] - g'(a^*_{imt}) = 0. \tag{2.8}$$

[30]Using the envelope theorem, we obtain the following relation: $\frac{\partial \pi_{ims}}{\partial a_{imt}} = \frac{\partial \pi_{ims}}{\partial w_{imt}} \frac{\partial w_{ims}}{\partial a_{imt}}$.

The effort in period t impacts the quality and grades in this period that influence all future profits. The driver equates the marginal benefit, which is the increase in future profits, with marginal cost, namely, the derivative of the cost of effort function. The higher the impact of the current effort on future profits is, the greater the effort. Proposition 5 characterizes optimal levels of effort throughout the driver's career.

Proposition 5. *The equilibrium sequence of effort tends asymptotically towards zero as the driver gains experience :* $\lim_{t \to +\infty} a_{imt} = 0$.

The proof of Proposition 5 is provided in Appendix 2.9. In the proof, we first use the consumer choice model presented above to rewrite drivers' profits. Second, we consider the limit of optimal effort as t goes to infinity (equation 2.8).

The key determinant of the driver's effort level is its impact on future profits. Initial reviews have a substantial influence on the posterior beliefs. Thus, the driver chooses a higher effort level. As more reviews become available, the residual uncertainty about the driver's type tends to zero, and the incentive to exert effort consequently tends to zero as well.

Proposition 5 generalizes the main result of Holmström (1999), by allowing elastic demand and introducing heterogeneity in the variance of types across populations. The trajectory of expected quality is illustrated in Figure 2.7.

Figure 2.7: Effort as a function of time

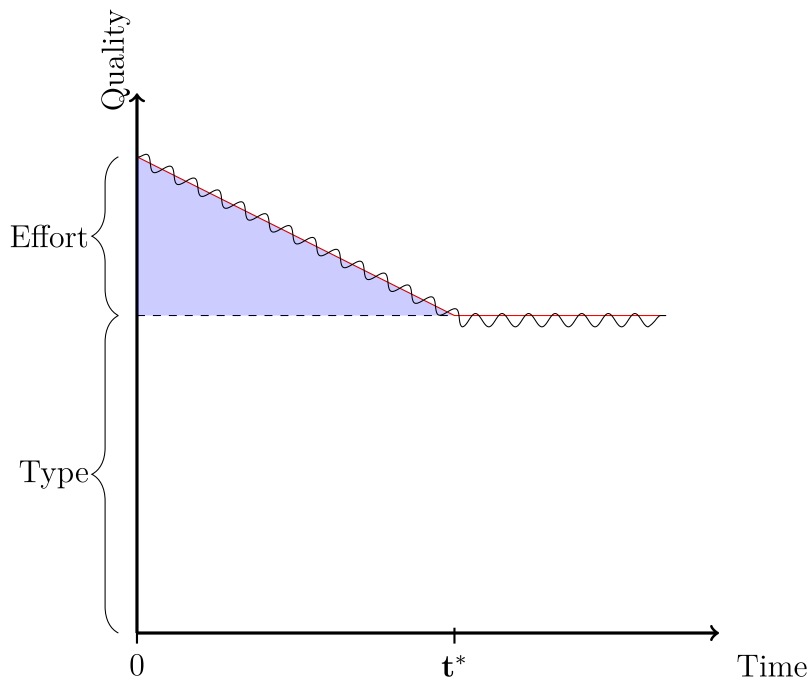

Considering the assumptions on the consumer choice, we can unpack the term $\mathbf{E}\left[\frac{\partial \pi_{ims}}{\partial a_{imt}}\right]$, which is the impact of effort in period t on profits in period s (see Appendix 2.9 for details). The profits of driver i from population m in period t are written as

$$\pi_{imt}(p_{imt}, X_{it}, w^{it}, c_i, \mathcal{S}_t) = M_t s_{imt}(p_{imt} - c_i), \tag{2.9}$$

where M_t is the number of passengers in period t, and s_{imt} is the probability of each of them purchasing seats from driver i; s_{imt} is given by equation 2.10,

$$s_{imt} = \frac{\exp(\alpha \mathbf{E}\left[w_{ijtm}|w^{it}\right] + \gamma p_{imt} + \beta r_{imt} + \mathbf{X_{imt}}\theta)}{1 + \sum_{k=1}^{N}(\exp(\alpha \mathbf{E}\left[w_{kjtm}|w^{kt}\right] + \gamma p_{kmt} + \beta r_{kmt} + \mathbf{X_{kmt}}\theta)} \tag{2.10}$$

Under our assumptions on the profit function and the choice rule of passengers, the impact of a unit of effort in period t on profits in period s is described by

$$\frac{\partial \mathbb{E}(\pi_{ims})}{\partial a_{imt}} = \frac{h_\epsilon}{h_{ims}} \frac{\alpha}{\gamma} \mathbb{E}[M_s s_{ims}]. \tag{2.11}$$

Two elements determine the optimal level of effort: the informativeness of a review ($\frac{h_\epsilon}{h_{ims}}$) and the impact of an increase in quality on market share ($\alpha/\gamma\mathbb{E}[M_s s_{ims}]$). The precision of the reputation system h_ϵ increases the level of effort. The more informative a review is, the higher the gain from exerting effort. The term $h_{ims} = h_m + sh_\epsilon$ reflects how much uncertainty remains about the driver's type; h_m is the inverse of the variance of types in population m: the higher the variance is, the greater the effort.

Next, the ratio of elasticity of demand with respect to quality α to the elasticity with respect to price γ impacts the optimal choice of effort. The more the market cares about quality, the greater are the efforts. Finally, higher future market shares $\mathbb{E}[M_s s_{is}]$ also increase the optimal level of effort.

The possible discrepancy between the driver and the market in the prior beliefs about the distribution of types affects efforts being made. Corollary 1 shows that drivers facing incorrect and overly pessimistic beliefs about their types exert greater efforts than they do in the case of the market and the driver agreeing on the lower belief.

Corollary 1. *A driver i from population m of mean type $\hat{\mu}_m$, facing overly pessimistic beliefs about the expected type $\mu_m < \hat{\mu}_m$, exerts greater effort than the level anticipated by the market:*

$$a_{imt}(\mu_m, \hat{\mu}_m) > a_{imt}(\hat{\mu}_m, \hat{\mu}_m),$$

where $a_{imt}(\mu_m, \hat{\mu}_m)$ stands for the optimal level of effort exerted in period t by driver i from population m with mean type $\hat{\mu}_m$, facing a market belief that the mean type in her population is μ_m.

The details of the proof are provided in Appendix 2.9. The optimal level of effort is characterized by equating the marginal return from providing effort with the marginal cost of effort. If drivers expect the reviewing process to improve the beliefs about the quality of their service, they will exert greater effort than expected by the market; this is so because their marginal return from effort is higher than what the market believes.

Pricing stage

On BlaBlaCar, drivers rarely compete against each other more than once. Therefore, the motive to deny the future advantage (in the form of a higher reputation) seems not to matter for current prices. By **Assumption 2**, a driver takes into account only the impact of current prices on future reputation and neglects the impact on the reputation of future competitors.[31]

Drivers choose prices and efforts to maximize future profits, which depend on the future market structures $\mathcal{S}_s$. It is unclear, a priori, what information BlaBlaCar drivers are taking into consideration when they form this expectation. By **Assumption 3**, we assume that a driver chooses optimal prices as if the future market structure were identical to the current one.

Two state variables influence optimal prices. The first is the market perception of the driver's quality $\mathbf{E}\left[w_{ijtm}|w^{it}\right]$. In the case of the market holding an incorrect prior belief, the driver knows that obtaining reviews on average improves the posterior belief. The second is the number of reviews r_{imt}. As a consequence, drivers choose prices taking into account the current market structure, the probability of transitioning into future states (obtaining a review), and the expected quality following each transition into a different state (reputation level). Therefore, the observed prices are solutions to the following problem:

$$\begin{aligned} p^*_{imt} = &\arg\max\{\pi_{imt}(p_{imt}, X_{imt}, w^{it}, c_i, \mathcal{S}_t) \\ &+ \delta\Sigma\left[\pi_{imt+1}(p_{imt+1}, X'_{imt+1}, w^{imt'}, c_i, \mathcal{S}_t)\right] \times p(X'_{imt+1}|X_{imt}) \times p(w^{imt'}|w^{imt}) + ...\}, \end{aligned} \tag{2.12}$$

where $p(w^{imt+1}|w^{imt})$ and $p(X'_{imt+1}|X_{imt})$ summarize transition probabilities,

[31] By **Assumption 2** we deviate from most of the empirical IO literature on dynamic competition, that typically takes into account the impact of today's behavior (pricing) on a future market structure (see, for example, Besanko et al. (2014); Doraszelski and Pakes (2007)).

and δ is a discount factor. $p(w^{imt+1}|w^{imt})$ summarizes the probability of obtaining each grade conditioned on having a reputation level w^{imt}. Intuitively, the higher the average grade in period t is, the higher the expected grade in period $t+1$. Higher grades lead to higher expectations of quality, and as a result, a higher expected utility of passengers. The latter transition probability captures the probability of obtaining a grade; it is a function of prices. Lower prices imply a higher chance of selling a seat and receiving a review.

Putting together the decisions of passengers and drivers, we obtain a definition of an equilibrium:

Definition 1. *An equilibrium in a ridesharing market is a set of*

1. *purchasing decisions of passengers that maximize their expected utility conditioned on their priors μ_m $\forall_m$, and*
2. *optimal price p^*_{imt} and effort: a^*_{imt} for each driver i characterized by marginal cost and type (c_i, η_i) such that*
   - $a^*_{imt} = g^{-1'}\left(\sum_{s=1}^{+\infty} \beta^{s-t} \times \mathbf{E}\left[\frac{\partial \pi_{ims}}{\partial a^*_{imt}}\right]\right)$ $\forall_i$, *and*
   - p^*_{imt} *is a solution of equation 2.12* $\forall_i$.

2.5 Identification and estimation

In this section, we present assumptions under which we can identify parameters of interest and estimate them. There are generally three groups of parameters. First, demand elasticities α, γ, β and θ. The key observables are prices and the numbers of sold seats. Additional available information is provided by the conditioning variables r and X for all drivers in each market. The second group consists of parameters related to the model of belief formation and updating, where we are interested in the prior beliefs μ_m, true distributions of types $\mathcal{N}(\hat{\mu}_m, 1/h_m)$ in each population m, and the informativeness of the reputation system h_ϵ. The observables that we will use to recover these parameters are the histories of grades of individual drivers and their market outcomes. Finally, the third group contains supply-side parameters including drivers' types η_i, efforts a_{imt}, and marginal costs c_i. We will also identify the cost of effort function $g(a_{imt})$. The observables of the supply-side involve the prices set by drivers, the histories of their grades, and transition probabilities $p(w_{imt+1}|w_{imt})$, $p(X'_{imt+1}|X_{imt})$.

2.5.1 Demand estimation

We propose a standard conditional logit model of demand. Here, we discuss some of its main features; detailed proofs and further discussion are provided in McFadden (1974). We assume that the utility of passengers is linear in the characteristics of drivers, that is,

$$u_{ijtm} = \alpha \mathbf{E}\left[w_{itm}|w^{it}\right] + \gamma p_{it} + \beta r_{it} + \mathbf{X_{it}}\theta + \varepsilon_{ijt},$$

where subscript j refers to passengers. In our baseline model, the stochastic term ε_{ijt} is the only difference between the passengers. We assume that it is a random variable with an extreme distribution $\mathcal{F}(\varepsilon_{ijt}) = \exp(-\exp(-\varepsilon_{ijt}))$. The probability that passenger j chooses driver i from N available drivers (indexed by k) and the outside option is $P_{ij} = P(u_{ij} \geq u_{ik}, \forall_{k \neq i})$, which given the assumption on u_{ijtm}, is

$$P_{ij} = \frac{\exp\left(\alpha \mathbf{E}\left[w_{itm}|w^{it}\right] + \gamma p_{it} + \beta r_{it} + \mathbf{X_{it}}\theta\right)}{1 + \sum_{k=1}^{N} \exp\left(\alpha \mathbf{E}\left[w_{ktm}|w^{kt}\right] + \gamma p_{kt} + \beta r_{kt} + \mathbf{X_{kt}}\theta\right)},$$

where the utility of the outside option is normalized to zero. McFadden (1974) shows that the log-likelihood function with these choice probabilities is globally concave in the parameters of demand. Thus, we can estimate its parameters by maximizing likelihood function with M observations (passengers),

$$\max_{\alpha,\gamma,\beta,\theta} \sum_{j=1}^{M} \sum_{i=1}^{N} d_{ijt} \ln P_{ij}(\alpha, \gamma, \beta, \theta)$$

,where $d_{itj} = 1$ if passenger j chooses driver i, and $d_{itj} = 0$ otherwise.

The identifying assumption is that our controlling variables X_{it} capture all demand-relevant driver-specific characteristics so that there is no heterogeneity across drivers that is observed by passengers but not by us. We make this assumption because, in our dataset, we indeed observe all information that is available to passengers. Nevertheless, for robustness, we introduce instrumental variables (cost shifters) in Appendix 2.9 to control for potential endogeneity, and also introduce random coefficients.

The number of potential passengers M is measured directly in our dataset. We have previously used the number of clicks that each listing received to measure the respective listing's popularity. However, the total number of clicks in the market can proxy the number of potential passengers. Within a market, defined as a route-and-day combination, we use the highest number of clicks received by any listing to represent the total number of passengers that have been interested in booking a ride. The difference between the maximum number of clicks and the total number of sold seats proxies the number of passengers that have searched for a ride, but did not buy. In other words, the latter passengers chose their outside option. The market size measured in this way exhibits significant time variation.

Market prior beliefs: We do not observe the market's belief about the expected quality $\mathbf{E}\left[w_{itm}|w^{it}\right]$. However, we know that passengers' beliefs converge to underlying quality as drivers receive reviews. Thus, drivers who have accumulated a substantial number of reviews face correct beliefs, which are consistent with their observed reputation. To recover market beliefs about drivers with no or few reviews, we will first estimate demand using a subset of markets where there are only experienced drivers (10 thousand out of 60 thousand markets).

In the second step, we use the estimated demand to predict the expected number of sold seats for the entire dataset. If, for a subset of drivers (for example, minority drivers), passengers at the booking stage are systematically incorrect about the grade they will give after the ride, the predicted market share obtained with our model will differ from the observed number of sold seats. We will use this difference to obtain the disparity between the grade given after the trip and the market expectation of a grade. To do that, we compare the market outcome s_{imt} to the prediction and assign the entire prediction error to passengers' errors in the assessment of the expected quality $\tilde{w}_{imt}$:

$$\tilde{s}_{imt} - s_{imt} \propto \tilde{w}_{ijt} - \mathbf{E}\left[w_{imt}|w^{imt}, \mu_m\right], \tag{2.13}$$

where $\mathbf{E}\left[w_{imt}|w^{imt}, \mu_m\right]$ is the market's belief about the expected quality of driver i, from population m with a history of grades w^{imt}.

Furthermore, from the model of belief formation and updating, we obtain a functional form of the expected quality. We attribute the difference to the disparity between the belief about the mean type in population μ_m and the actual mean $\hat{\mu}_m$.

2.5.2 Supply-side parameters

The key supply-side parameters are drivers' types η_i, their efforts a_{imt}, and marginal costs c_i. For all drivers in our dataset, we have histories of ratings obtained from the driver's first ride until the moment the driver appears in our dataset for the last time. We will use these grades to recover drivers' types and efforts.

Figure 2.3 of section 2.3.2 shows the average ratings at different stages of drivers' careers.[32] We observe that the ratings are high in the beginning and stabilize as more reviews become available. The observed trajectory of grades is consistent with the prediction of the model - the initial increase in grades is due to efforts, while the level at which the grades stabilize coincides with the driver's type.

[32] The first point on the left chart is the average first grade. We restrict the sample to drivers who stayed on the platform long enough so that they gathered enough reviews to reveal their types. Restricting the sample has an additional advantage of mitigating the survivorship bias stemming from the selection of the drivers with high grades. As pointed out in section 2.3, receiving a low grade increases the chance of a driver leaving the platform; thus, the grades of drivers who stayed were on average higher than the ratings of those who left the platform early on.

By Proposition 5, the optimal level of effort approaches zero as t tends to infinity. We assume a burnout period t^*, after which the level of effort is low.[33] Thus, we define the parameters of interest as follows:

- The intrinsic quality (type) of an individual driver is the average of her grades after t^*,

$$\eta_i = \frac{\sum_{t=t^*}^{T} w_{imt}}{T - t^*}, \tag{2.14}$$

  where T is the last period in which we observed a grade given to driver i.

- The effort a^*_{imt} of driver i from population m at time t with history of grades w^{it} is

$$a^*_{imt} = \frac{\sum_{s=1}^{N^{m,w^{st}}} (w_{smt} - \eta_s)}{N^{m,w^{st}}}, \tag{2.15}$$

  where s indexes drivers from population m with history of grades w^{st}, and $N^{m,w^{st}}$ is the number of such drivers in our dataset. Thus, the expected effort of driver i is the average difference between grades and types for all drivers with the same characteristics (including the number of reviews), types, and histories of grades.

- We assume that the distribution of the error term is normal with zero mean. We are interested in estimating the precision (the inverse of variance) of the error term, which is given by the inverse of the mean of variances of grades after t^*,

$$h_\epsilon = \frac{N^{t^*}}{\sum_{s=1}^{N} Var(w_{st})} \forall_{t>t^*}, \tag{2.16}$$

  where N^{t^*} is the set of grades of drivers with $t > t^*$.

[33] In practice when it is no longer statistically significant for both minority and nonminority drivers.

We need several assumptions to identify these parameters in the data. First, there are no listing-specific variables other than types, efforts and exogenous errors that influence grades. In particular, we assume that prices do not influence grades. Appendix 2.9 provides some evidence supporting this assumption. Second, error terms are random variables, with mean zero. We require that: $\mathbf{E}[\epsilon_{it} + \epsilon_{it+1}] = \mathbf{E}[\epsilon_{it}] + \mathbf{E}[\epsilon_{it+1}]$. This is necessary, so that

$$\lim_{T\to\infty}\left[\frac{1}{T-t^*}\sum_{t=t^*}^{T}(\eta_i + \epsilon_{it})\right] = \eta_i.$$

Next, in order to identify the optimal level of effort, we need that the error term is independent across drivers and that there are no unobserved listing-specific and demand-relevant characteristics that influence future market shares so that:

$$a^*_{imt} = \mathbf{E}\left[a^*_{kt} + \epsilon_{kt} | X_{kt} = X_{it}, w^{kt} = w^{it}, \eta_i = \eta_k\right] = \lim_{N^{m,w^{st}}\to\infty}\left[\frac{1}{N^{m,w^{st}}}\sum_{s=1}^{N^{m,w^{st}}}(w_{smt} - \eta_s)\right],$$

where $N^{m,w^{st}}$ is the subset of drivers that have the same incentive to exert effort as driver i. In this way, we argue that drivers with the same observed characteristics $X_{kt} = X_{it}$, same type $\eta_k = \eta_i$ and the same history of grades $w^{kt} = w^{it}$ exert the same level of effort. To be able to identify the optimal effort for all drivers, we rely on a large number of observations, so that in each type, characteristics, and grades combination, there are enough drivers.

Having determined the types of individual drivers, we obtain the distributions of types in different populations. The mean is given by

$$\hat{\mu}_m = \frac{1}{N^m}\sum_{s=1}^{N^m}\eta_s,$$

and the precision is

$$h_m = \frac{N^m}{\sum_{s=1}^{N^m}(w_{smt} - \eta_s)^2},$$

where N^m is the number of drivers in population m.

Estimation of the cost of effort function: The cost of effort function is unknown, we have assumed that it is convex and increasing. The function $g(\cdot)$ defines the optimal level of effort by equating the marginal benefit from exerting a unit of effort with the cost of such a unit.

$$a_{imt} = g^{-1'}\left(\sum_{s=t}^{n} \beta^{s-t} \frac{h_\epsilon}{h_{mk}} \frac{\alpha}{\gamma} \mathbb{E}[M_k s_{imk}]\right). \tag{2.17}$$

From the discussion above, we know how to measure the levels of effort. The arguments of the function have also been already identified. The elasticities of demand with respect to quality and price α and γ are estimated with the demand model. We recover the variance of types and that of the error term from the observed grades. We can obtain the future market shares by predicting the demand for driver i in future periods. Thus, we can approximate function $g(\cdot)$.

Appendix 2.9 shows estimates of polynomials of various degrees and compares the fit of the model. In the baseline case, we will assume that $g(\cdot)$ is a quadratic function.

Pricing stage

The marginal cost of having a passenger on board defines the profitability of using the platform. We argue that drivers act strategically while setting their introductory prices, and in section 2.3.2, we provide some evidence of this. In the reduced-form results, we also show that the returns from reputation are decreasing. There is a saturation point at approximately ten reviews, after which there are no more incentives to reduce prices to receive more reviews. Thus, prices in periods from the first to the tenth exhibit the static-dynamic tradeoff, while prices in periods after the tenth can be interpreted as static profit-maximizing prices.

Assuming that after the tenth period, prices maximize static profits, we can recover marginal costs. To this end, we need to first estimate markups, which we obtain from estimated demand. The difference between the price and the markup for prices of drivers with more than ten reviews is given by

$$p^*_{imt} = c_i + \frac{s_{imt}}{\frac{\partial s_{it}}{\partial p_{imt}}}. \tag{2.18}$$

In this way, we recover marginal costs and obtain their distribution for drivers who stayed at least until the tenth period.

As we argued while discussing the model, each price is a solution of a dynamic programming problem. We observe the transition rules $p(w^{imt+1}|w^{imt})$ and $p(X'_{imt+1}|X_{imt})$ directly in the data; $p(w^{imt+1}|w^{imt})$ is the probability of obtaining each possible grade conditioned on having the level of reputation w^{imt}, and $p(X'_{imt+1}|X_{imt})$ is the probability of receiving a grade after selling a seat.

To find the optimal price in period t, we need to first characterize the optimal behavior in period $t+1$, because the value of being in period $t+1$ defines the incentive to get there. Hence, to solve the problem we will proceed by backward induction. First, in period ten, we assume Bertrand pricing. We find the optimal price for a driver with a given set of characteristics and a marginal cost. We also obtain the value of being in period ten (the discounted sum of profits). Then, in period nine, there is already an incentive to reduce the price to proceed faster to period ten. Hence, the problem in period nine is written as

$$\begin{aligned} p^*_{im9} = \arg\max\{&(p_{im9} - c_i)M_9 s_{im9}(p_{im9}, X_{im9}) + \delta\left[s_{im9}V_{im}(10)\right. \\ &+(1-s_{im9})\left((p_{im9}-c_i)M_9 s_{im9}+\delta(s_{im9}V_{im}(10)+(1-s_{im9})(p_{im9}-c_i)M_9 s_{im9}+...)\right)]\}, \end{aligned} \tag{2.19}$$

where $s_{im9}(p_{im9}, X_{im9})$ is the probability of selling a seat given the price p_{im9} and characteristics X_{im9}, $M_9 s_{im9}(p_{im9}, X_{im9})$ is the expected number of sold seats, which determines the expected number of new reviews, and $V_{im}(10)$ is the expected value of being in period ten (expected with respect to the grade that i will obtain). If driver j does not sell a seat, she solves the problem of period nine again until she obtains a review.

After determining the optimal price for period nine, we proceed to period eight and so forth until we reach the task of determining introductory prices.

Any price that we observe is the solution to this problem; thus, we can identify the marginal cost from each observed price.

2.6 Results

In this section, we present and discuss the results of the estimation of the model. First, we show estimates of demand. Second, we demonstrate how incentives to exert effort differ across minority and nonminority drivers. Next, we present the estimated prior. Finally, we show estimates of marginal costs and discuss how the incentives to invest in reputation depend on them.

Demand estimates

In Table 2.6, we present the results of demand estimation. The dependent variable is d_{ijt}, a binary variable that takes the value one if driver i was selected by passenger j and is zero otherwise. As we proceed from column one to column two, we add more controls. The variable *type* is the average grade from the tenth onwards, while *reputation* (column three) takes into account all grades available on drivers' profiles. The regression presented in column one controls for the type, the number of reviews, and price; in column two, we add a full set available controls, time, and trip specific effects. In column three, we use *reputation* instead of *type*. Demand is estimated using a subset of 10241 markets (400 thousand choice situations). We will use model two in the supply-side estimation and the analysis of counterfactuals.

Demand is generally not very elastic. The elasticity with respect to price is -0.12 and 0.57 with respect to the expected quality.

Market prior beliefs

To recover passengers' beliefs about the expected quality of service offered by drivers with no reputation, we first predict the number of sold seats using the *model 2* fromTable 2.6. Next, we attribute the error of the prediction to the expectation of the grade.

Table 2.6: Demand estimates

	Model 1	Model 2	Model 3
Ride price	-0.00 (0.00)***	-0.00 (0.00)***	-0.00 (0.00)***
Type	0.12 (0.06)**	0.13 (0.06)**	
Log(number reviews)	0.15 (0.02)***	0.14 (0.02)***	0.15 (0.02)***
Minority		0.06 (0.05)	
Reputation			0.24 (0.10)**
AIC	31929.66	30150.03	30145.51
R^2	0.45	0.45	0.45
Max. R^2	0.49	0.49	0.49
Num. events	154259	147905	147905
Num. obs.	470165	442839	442839

Note: Demand estimates: subset of markets. All coefficients presented in the Appendix 2.9 For additional robustness check we also estimate a model with random coefficient associated with price (BLP); elasticities of price and quality are in the Appendix 2.9.

We find that minority drivers with no reputation are expected to deliver the quality of 4.16 (on a scale of 1 to 5), while they are graded at 4.619 on average. The expectation corresponds to the 7.5th percentile of the distribution of quality. For comparison, nonminorities are expected to provide a quality of 4.59 and are graded 4.68. Figure 2.8 summarizes this. The solid blue line represents the average first grade obtained by minority drivers, while the solid red line is the market expectation of the grade. Dotted lines correspond to grades and their expectations for nonminority drivers. The distribution of all grades is shown in black.

Finally, as argued throughout this paper, the beliefs about quality are being updated; thus, the two numbers converge. Minority non-entrants (with more than two reviews) are believed to be of quality 4.539 before the trip and are graded 4.592 ex post.

Figure 2.8: Erroneous beliefs and given grades

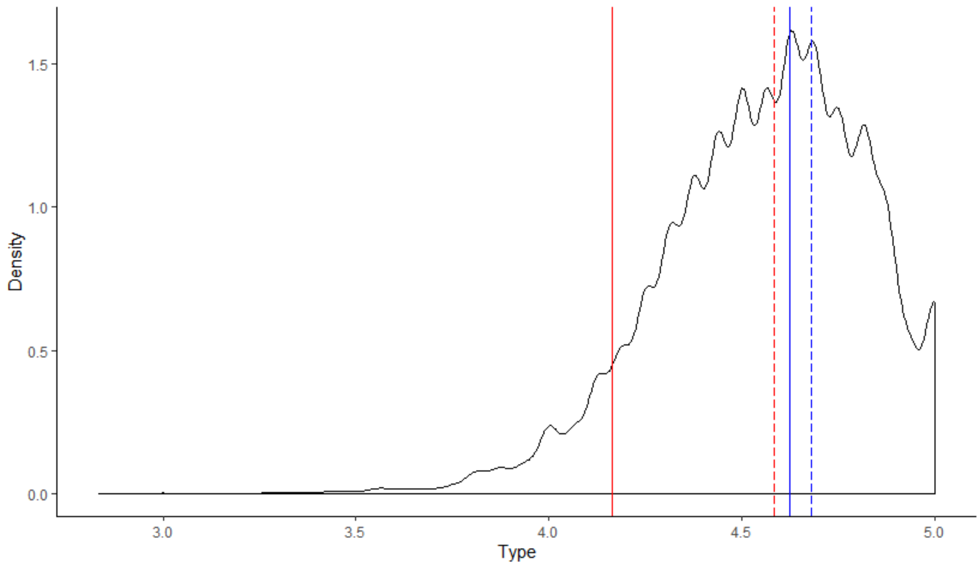

Note: The distribution of grades is shown in black. Blue lines represent the mean first grade obtained by a minority driver (solid line) and by nonminority drivers (dotted line). Red line illustrate the market beliefs on the expected quality (minority- solid line, nonminority - dotted line).

The expected quality in the first period depends only on the prior belief about the distribution of quality among minority drivers and the expected level of effort. Consistently with the result of Corollary 1, minority drivers exert greater effort than the market expects them to. Given the estimated parameters, the difference between the two levels of effort results in the difference of 0.04 in the first grade.

The distribution of types in the population of nonminorities has mean of 4.56 and variance 0.07. For the population of minority drivers, the mean is 4.49, and the variance is 0.09. We account for the difference in expected and exerted efforts and find that the market expects the mean type of minority drivers to be at 4.09.

Incentives to exert effort

The incentive to exert effort is determined by the magnitude of the impact of a unit of effort on future profits. Figure 2.9 shows the average (across all drivers) increases in the next period's profits due to a unit of effort exerted in the current period, determined as $\frac{h_\epsilon}{h_{mk}}\frac{\alpha}{\gamma}\mathbb{E}[M_k s_{ik}]$. We show how this quantity changes during a drivers' career. The expected market shares are those observed in the data; elasticities of demand are from model three in Table 2.6. Red dots indicate the return to effort for minority drivers and blue dots represent the corresponding results for nonminority drivers.

First, the impact of efforts on profits decreases as more information about drivers becomes already available. Second, the initial increase is higher for minority drivers. Two countervailing factors shape the disparity between minority and nonminority drivers: a higher variance of types in the population of minority drivers results in more uncertainty about individual types, and as a consequence, higher efforts. Although the expected profits in the first several rounds are smaller for minority drivers, which dampens the incentive to exert effort, the market shares increase over time, so that the latter effect is not particularly strong.

The incentives to exert effort are closely linked with the impact of a grade on future revenue. Table 2.7 shows the change in revenue following a grade from one to five. We take into account the elasticity of demand with respect to the number of reviews and quality. Only a grade of five has a positive impact. The grade of four leads to almost no change in revenue, and all lower grades result in negative and substantial changes. Minority drivers experience a more significant reaction to any grade because of the higher variance of types. They lose more as a result of low grades and experience a more significant benefit from a grade of five.

Marginal costs

Figure 2.10 shows the distribution of the recovered marginal costs on selected trips. These costs are related to trip length; long trips are associated with

Figure 2.9: Incentives to exert effort

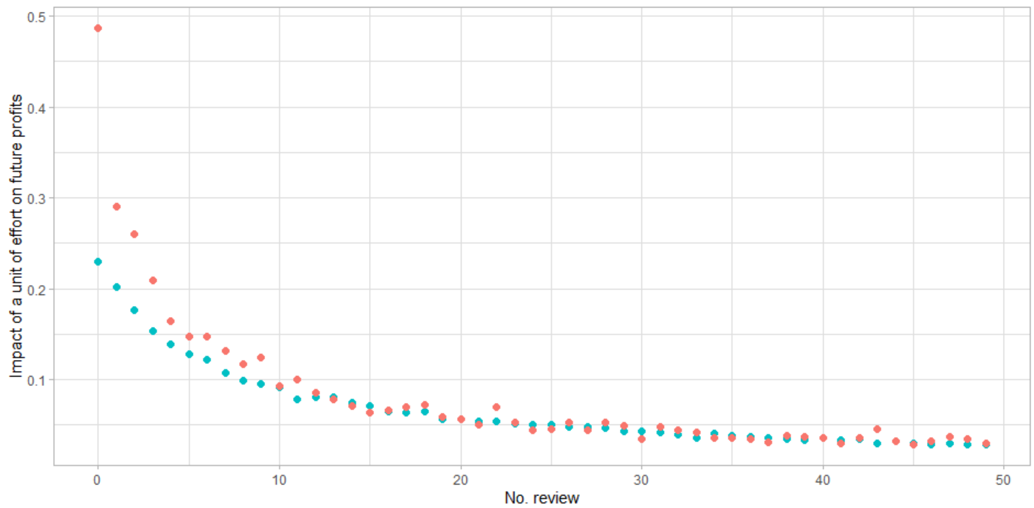

Note: Horizonal axis - number of reviews. Vertical axis - the impact of a unit of effort on future market shares. Red dots - minority drivers. Blue dots - nonminority drivers.

Table 2.7: Impact of a grade on revenue

	1	2	3	4	5
Minority	-33.18%	-21.98%	-10.78%	-0.42%	11.63%
Nonminority	-45.76%	-29.75%	-11.74%	-5.28%	22.28%

Note: The figures show percentage changes of the next predicted revenue amount following a grade from 1 to 5. This impact arises from the the number of reviews and expected quality.

higher marginal costs than are shorter trips. The difference in marginal costs between minority and nonminority drivers (23.3 and 22.6, respectively) is 3.2%.

Figure 2.10: Marginal costs

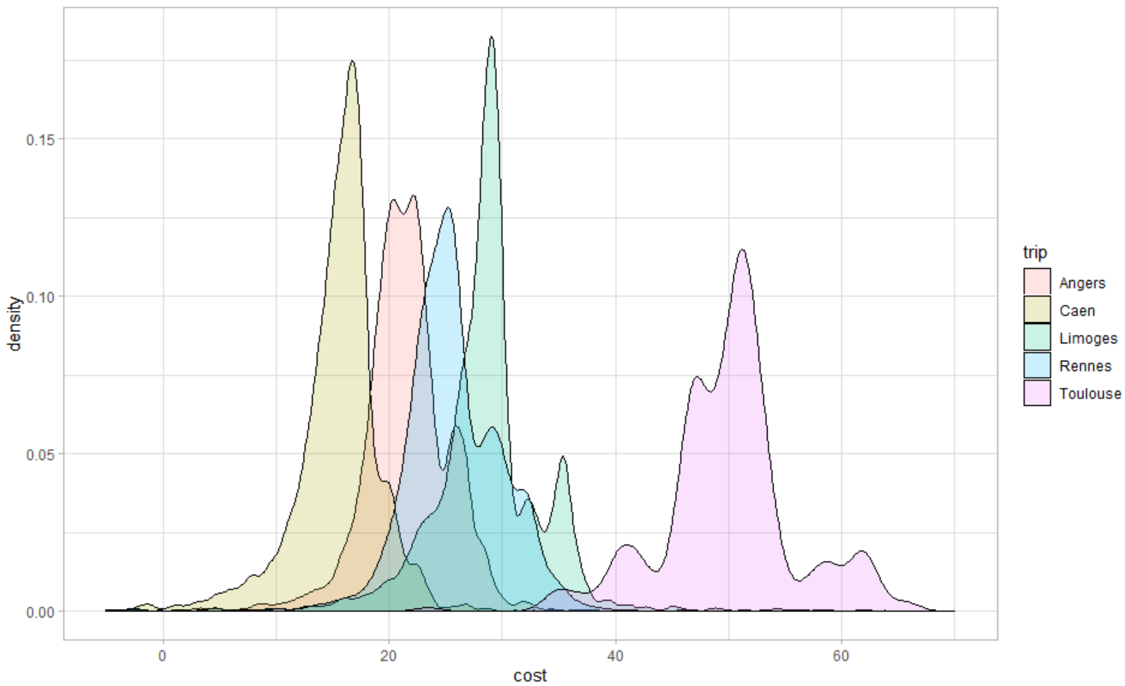

Note: Marginal costs given by equation 2.18 of drivers with more than ten reviews; selected trips

Pricing results

We are interested in how the incentive to invest in reputation translates into low introductory prices. Throughout this section, we will compare the prices set by drivers if they internalize the reputation-building incentive while setting prices and the prices set if drivers do not do so. Dynamic prices are solutions of equation 2.12, while static prices satisfy equation 2.18.

Figure 2.11 compares static and dynamic prices of minority and nonminority drivers. We fix the marginal cost and all other driver-listing-specific characteristics, except for the number of reviews and the expected quality.

The expected quality in each period equals the mean of expected qualities of all drivers with the same set of characteristics.[34]

Optimal static prices (bullets in Figure 2.11) increase over time due to the positive elasticity of demand with respect to the number of reviews. Nonminority drivers receive, on average, the same reviews as the market expects. However, minority drivers experience an additional benefit from reputation because of the increase in posterior beliefs. Thus, even static prices increase more rapidly for minorities than for nonminorities.

The prospect of higher profits motivates all drivers to act strategically and offer discounts. Dynamic prices start at lower levels (e.g., the first period's prices might be below costs), increase more rapidly, and converge to static prices at period ten. Note that under dynamic pricing, drivers sell seats faster during the first couple of periods. Minority drivers take into account the expected correction in the market belief about their quality. As a result, they offer larger discounts.

The change in prices from period to period also depends on marginal costs. Figure 2.12 repeats the exercise illustrated by Figure 2.11 but considers several levels of costs. We present the difference in introductory prices between static and dynamic pricing modes for different levels of costs. The difference increases with marginal costs. Lowering introductory prices increases the probability of selling a seat and receiving a review. However, at lower levels of marginal costs, the prices are already relatively low in the static case. Thus, lowering them further has a proportionally smaller impact on increasing the chance of receiving a review. Furthermore, drivers with low marginal costs earn a significant markup even when they have only a few reviews. Thus, their incentive to invest in reputation is smaller.

So far, we have focused on one driver with one set of characteristics. However, we have already observed that the incentive to offer a discount from static

[34]The demand predicted for a ride with: a photo, the automatic acceptance feature, the maximum 2 passengers option, the ride occuring during the day, the time since the listing has been posted equal to the mean in the dataset, the notice equal to the mean in the dataset, seniority equal to the mean in the dataset, car price equals to the mean in the dataset, the ride occuring during a weekday on a non-strike day.

Figure 2.11: Dynamic vs. static prices

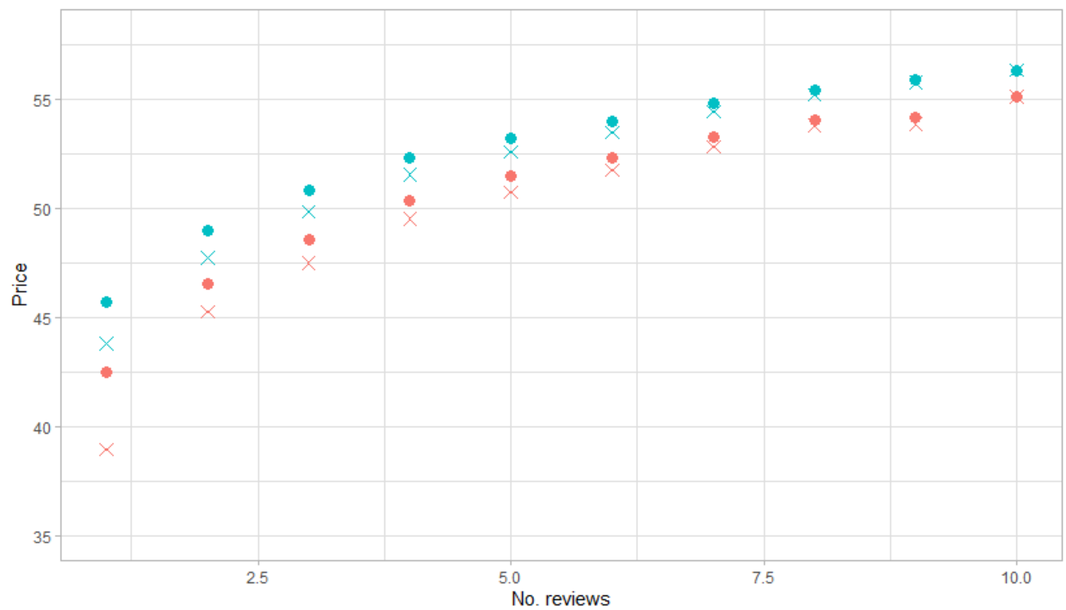

Note: Horizontal axis - the number of reviews. Vertical axis- the optimal price. Minority drivers - red; nonminority - blue. Bullets- static profit-maximizing prices. Crosses- prices, resulting from internalizing reputation-building incentive.

Figure 2.12: Discount in introductory prices for various levels of marginal cost

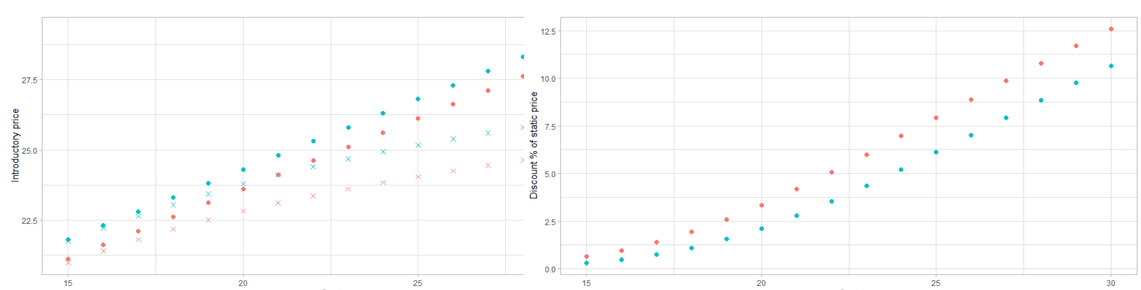

Note: Static vs. dynamic introductory prices. Minority drivers - red, nonminority drivers - blue. The left panel compares dynamic prices (crossess) with static prices (bullets). The right panel shows the difference between static and dynamic as a percentage.

prices for drivers with few (or no) reviews depends on the level of marginal costs. It also depends on other drivers' characteristics.

Generally, the lower the initial market share is, the higher the relative increase in profits following a review. Therefore, to quantify the average discount in introductory prices, we set the parameters in the algorithm to match those of listings we have observed (thus, we assume that ride-specific parameters - photograph, automatic acceptance, weekday, etc. do not change as the driver receives reviews), and estimate marginal costs. We focus on the sample of markets used in demand estimation. Based on the recovered marginal costs, we compute introductory prices for a driver who follows Bertrand pricing and compare them with the observed dynamic prices.

We find that nonminority drivers reduce their prices by 4.08% on average, which is a significant investment in reputation. Minority drivers reduce prices by 8.03%; the larger discount is due to a higher increase in future market shares following an expected review. Consistently with the example in Figure 2.12, the difference is higher for drivers with higher marginal costs.

2.7 Counterfactual experiments

The structural model allows us to generate counterfactual experiments. We will analyze three alternative scenarios. First, we simulate market outcomes under the correct prior. In this scenario, passengers have correct beliefs about the expected quality of all minority drivers. Comparing the baseline scenario with this experiment allows us to calculate the cost to minority drivers of erroneous beliefs. Second, we study the market in which the gap between minority and nonminority drivers remains constant. In this case, the expected quality is always reduced by the size of the bias. This simulation highlights the difference between statistical discrimination (the baseline case) and taste-based discrimination. Finally, we evaluate a policy intervention proposed by Benjamin Edelman and Michael Luca (Edelman et al. (2017)) that makes the profiles of drivers ethnicity-blind. Table 2.8 summarizes the main results.[35]

Table 2.8: Summary of counterfactuals

	Δ quality	Δ efforts	Δ p_1 minority	Δ π minority	Δ π n.minority
Correct prior	2.9%	4.91%	3.91%	19.13%	-0.48%
Persistent bias	-4.95%	-14.28%	0.62%	-7.69%	0.11%
Debias	10.03%	7.54%	13.34%	21.6%	-0.85%

Note: All values are percentage changes compared to the baseline. Column 1: average change in the expected quality of minority drivers on trips 1-15. Column 2: change in total efforts of minority drivers. Column 3: change in introductory price charged by minority drivers. Column 4: change in average profits of minority drivers over trips 1-15. Column 5: change in average profits of nonminority drivers over trips 1-15.

Cost of the incorrect prior

This exercise aims to quantify the cost of erroneous beliefs. Under this scenario, minority drivers will be evaluated ex ante following their true quality, as revealed by the grades they obtain ex post.

This change spurs several reactions. First, minority drivers will be perceived by the market as being of higher quality. They will be able to raise prices and exert more effort, so their quality will increase further. Nonminority drivers will react to this by reoptimizing their prices. Finally, passengers in the counterfactual markets will choose between minority drivers, whom they now perceive to be of higher quality, but who now charge higher prices, nonminority drivers with new levels of prices, and the outside option, that is unchanged.

This scenario assumes that the belief about the expected type of a minority driver with no reviews improves from 4.1 to 4.49 (a change from the 7.5th per-

[35] In each of the scenarios, we characterize a new equilibrium described by definition 1. Each of the proposed counterfactuals involves changes in passenger decisions, which leads to new optimal prices and efforts by both minority and nonminority drivers, which again lead to a different set of passenger decisions. Thus, we are looking for new vectors of purchasing decisions, pricing, and efforts such that none of the parties can gain by deviating.

centile to the 50th). The process of updating beliefs about individual drivers' quality proceeds the same way as before. As a consequence, throughout the first 15 periods, the average perception of the expected quality of minority drivers increases by 2.9%. Moreover, the amount of effort increases by 4.91%, further boosting quality. The optimal level of effort is particularly susceptible to changes in expected profits in the first period; hence, the sizable change.

The higher expected quality allows minority drivers to increase introductory prices. The incentive to reduce the price to hasten belief correction disappears. Introductory prices rise by 3.91%. A higher expected quality and the change in prices have a substantial effect on expected profits that increase by 19.13%.

In other words, 19.13% of profits is the price minority drivers have to pay for incorrect beliefs held by passengers. Finally, the profits of nonminority drivers decline by 0.48%. Most of the change in substitution is with respect to the outside option.

Persistent bias

Suppose that the bias against minority drivers is not subject to change. Each driver can have her individual reputation, but minority drivers are always considered to be worse, regardless of how many reviews they have. Let the size of this bias be given by the extent to which the expected belief about the type of minority drivers differs from the mean type revealed by the grades (4.1 vs. 4.49).

If there is no possibility of mitigating discrimination, minority drivers always achieve lower profits, and their incentives to exert effort vanish. The exerted efforts decline by 14.28%, further depressing the quality of service provided by minority drivers.

Interestingly, in this case, minority drivers will charge higher introductory prices. This is so because the expected quality of service of a minority driver with no reviews is the same as in the baseline case, but the incentive to reduce the price to receive more reviews is lower. The average profit throughout the first 15 periods is lower by 7.69%. There is little substitution away from minority drivers to nonminority drivers, whose profits increase by 0.11%.

Ethnicity-blind profiles

The publication of the studies by Edelman and Luca (2014) and Edelman et al. (2017) that were among the first to document racial bias on sharing economy platforms spurred a heated discussion about ways to address the problem. One of the proposals was to change the way platforms displayed ethnicity or gender-related information. To this end, the authors of the above papers developed a web browser plugin called *Debias Yourself*[36] that removed names and photos of hosts on Airbnb.

Airbnb itself started addressing the issue of racial bias by changing the way profiles were presented. In 2016, the listing page (the page displayed after a search query) stopped showing names and photos of hosts. Only information specific to the listing became available. To view host-specific information, a potential guest had to click on the listing.[37]

For this experiment, let us suppose that a passenger does not know whether the driver with whom she is planning on booking a trip is a minority. It is also impossible to deduce that from other observables. Therefore, the passenger forms an expectation based on the distribution of drivers in a given market. The share of minority drivers differs depending on the route. The highest ratio of minority to nonminority drivers is on the route from Lyon to Paris (16%), and the lowest is on the Rennes-Paris connection (7%).[38]

As a result, the market perceives minority drivers to have an expected quality that is higher by 10.03%. Drivers will react to this policy by reoptimizing effort levels and their pricing strategies. Now, both minority and nonminority drivers have an incentive to reduce their introductory prices because reviews improve the beliefs about quality for everyone. However, for minority drivers, this incentive is lower than in the baseline case. From a

[36]http://www.debiasyourself.org/index.html

[37]See https://www.cnbc.com/2017/04/07/airbnb-experimenting-with-site-design-to-fight-discrimination.html for details.

[38]We also assume that reviews do not reveal ethnicity.

static perspective, minorities should increase prices immediately because their quality is now believed to be higher. Considering both effects, we observe that introductory prices set by minority drivers increase by 13.34%. By the same logic, nonminority drivers reduce their prices.

The increase in the expected quality and the rise in prices result in higher profits of minority drivers; the latter increase by 21.6%. Nonminority drivers earn slightly lower profits, a reduction of 0.85%. The change in expected profits of nonminority drivers is more substantial for drivers with high marginal costs. In Figure 2.13, we show the change in percentage terms of the discounted sum of profits. This experiment reveals that if drivers are heterogeneous in unobservables, imposing a veil of ignorance on some observables might have unintended consequences.

Figure 2.13: Percentage change of the discounted sum of profits.

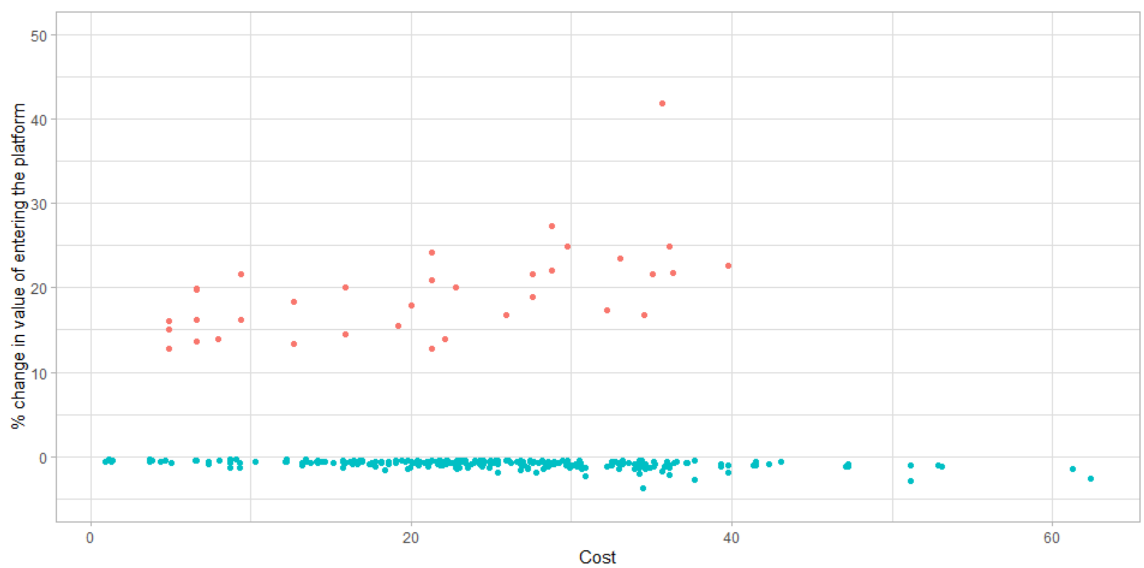

Note: Minority drivers - red. Nonminority drivers - blue. Results are for randomly selected 500 drivers. Horizontal axis- marginal cost. Vetical axis- change in the discounted sum of profits earned in the counterfactual scenario.

In this paper, we do not model entry into the market. However, given the changes in expected sums of profits in all three counterfactual scenarios, we should expect a change in the composition of drivers. Minority drivers have stronger incentives to join the platform when their expected quality is believed to be higher and under ethnicity-blind profiles. They would be less likely to enter the market when they face a persistent bias. The incentives to enter the market for nonminority drivers are changing in precisely the opposite direction.

2.8 Conclusions

Online discrimination against minorities has been documented in many prominent marketplaces. In this paper, we show that in the context of BlaBlaCar, a significant part of discrimination arises due to incorrect and overly pessimistic prior beliefs about the quality of service offered by minority drivers. These beliefs are altered with reviews. The initial gap of approximately 12% in revenue declines as minority drivers accrue reputation. The revenue differential for experienced drivers is statistically insignificant. The improvement in the performance of minority drivers is due to a causal effect of reviews, as we show using a difference-in-differences analysis.

This paper provides evidence that minority drivers use the reputation system to their benefit. They increase their levels of effort to receive high grades and set low introductory prices to build up their reputations faster. In the context of BlaBlaCar, the online reputation system allows mitigating ethnic discrimination. However, this is a costly fight for minority drivers. They have to persevere through an initial period of low economic outcomes and invest in their reputation. To calculate the cost of incorrect beliefs, we perform a counterfactual experiment. We simulate market outcomes in a scenario in which the initial beliefs about the quality of service of minority drivers are correct. Over the first 15 rides, we observe an increase in profits by 19%, which is the true cost of incorrect beliefs.

We propose a model of career concerns that represents a novel approach to studying the incentives of sellers in online markets. A reputation system creates an intertemporal externality. Reports of past performance can reveal some demand-relevant and seller-specific information and, as a result, boost or hurt future outcomes. However, the seeming randomness of reviews makes the task of extracting information out of grades difficult; this is why we need a model. We indeed show that reviews exhibit random components. Nevertheless, the model we propose allows us to separate the random element from the information that the market can use to update beliefs about the expected future quality. The ratings of minority drivers are on average higher than the market expectation. Thus, such drivers' quality of service is typically believed to be higher a posteriori than a priori.

The platform itself does not create prejudice against minorities. However, platform design can both mitigate discriminatory behavior and exacerbate it. BlaBlaCar provides information that reveals ethnicity of drivers, which allows passengers to discriminate based on it. The platform also equips minority drivers with tools to counter discrimination. The entire history of reviews is available on profiles of drivers, which helps inform future passengers. Drivers can also influence the speed of beliefs' updating by offering discounted prices. Thus, BlaBlaCar's online infrastructure enables its users to alter their incorrect priors.

This paper contributes to a long-standing discussion of the sources of discrimination. In our context, discrimination is to a large extent due to incomplete information. Passengers on BlaBlaCar are willing to change their beliefs when they are presented with an additional review. This result has clear policy implications: the provision of information is an effective way of tackling discrimination, at least in this market.

Bibliography

ABADIE, A. AND G. W. IMBENS (2016): "Matching on the estimated propensity score," *Econometrica*, 84, 781–807.

ABRAHAO, B., P. PARIGI, A. GUPTA, AND K. S. COOK (2017): "Reputation offsets trust judgments based on social biases among Airbnb users," *Proceedings of the National Academy of Sciences*, 114, 9848–9853.

AGRAWAL, A., N. LACETERA, AND E. LYONS (2016): "Does standardized information in online markets disproportionately benefit job applicants from less developed countries?" *Journal of international Economics*, 103, 1–12.

ALESINA, A., S. MICHALOPOULOS, AND E. PAPAIOANNOU (2016): "Ethnic inequality," *Journal of Political Economy*, 124, 428–488.

ALTONJI, J. G. AND C. R. PIERRET (2001): "Employer learning and statistical discrimination," *The Quarterly Journal of Economics*, 116, 313–350.

ARROW, K. J. (1973): "The theory of discrimination, discrimination in labor markets," *Achenfelter, A. Ress (eds.), Princeton–New Jersey.*

BANERJEE, A. AND K. MUNSHI (2004): "How efficiently is capital allocated? Evidence from the knitted garment industry in Tirupur," *The Review of Economic Studies*, 71, 19–42.

BARTOŠ, V., M. BAUER, J. CHYTILOVÁ, AND F. MATĚJKA (2016): "Attention discrimination: Theory and field experiments with monitoring information acquisition," *American Economic Review*, 106, 1437–75.

BECKER, G. (1971): "The Economics of Discrimination," *University of Chicago Press Economics Books.*

BESANKO, D., U. DORASZELSKI, AND Y. KRYUKOV (2014): "The economics of predation: What drives pricing when there is learning-by-doing?" *American Economic Review*, 104, 868–97.

Bohren, J. A., K. Haggag, A. Imas, and D. G. Pope (2019a): "Inaccurate Statistical Discrimination," Tech. rep., National Bureau of Economic Research.

Bohren, J. A., A. Imas, and M. Rosenberg (2019b): "The dynamics of discrimination: Theory and evidence," *American Economic Review*, 109, 3395–3436.

Carol, S., D. Eich, M. Keller, F. Steiner, and K. Storz (2019): "Who can ride along? Discrimination in a German carpooling market," *Population, Space and Place*.

Chiappori, P.-A., B. Salanie, and J. Valentin (1999): "Early starters versus late beginners," *Journal of Political Economy*, 107, 731–760.

Coate, S. and G. C. Loury (1993): "Will Affirmative-Action Policies Eliminate Negative Stereotypes ?" *The American Economic Review*, 83, 1220–1240.

Conlon, C. and J. Gortmaker (2019): "Best practices for differentiated products demand estimation with pyblp," Tech. rep., Working paper. url: https://chrisconlon. github. io/site/pyblp. pdf.

Cui, R., J. Li, and D. J. Zhang (2019): "Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb," *Management Science*.

DeGroot, M. H. (2005): *Optimal statistical decisions*, vol. 82, John Wiley & Sons.

Doraszelski, U. and A. Pakes (2007): "A framework for applied dynamic analysis in IO," *Handbook of industrial organization*, 3, 1887–1966.

Edelman, B., M. Luca, and D. Svirsky (2017): "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment," *American Economic Journal: Applied Economics*, 9, 1–22.

Edelman, B. G. and M. Luca (2014): "Digital Discrimination: The Case of Airbnb.com," *SSRN Electronic Journal*.

Farajallah, M., R. G. Hammond, and T. Pénard (2019): "What drives pricing behavior in Peer-to-Peer markets? Evidence from the carsharing platform blablacar," *Information Economics and Policy*.

Gandhi, A., Z. Lu, and X. Shi (2013): "Estimating demand for differentiated products with error in market shares," Tech. rep., cemmap working paper, Centre for Microdata Methods and Practice.

Ge, Y., C. Knittel, D. MacKenzie, and S. Zoepf (2016): "Racial and Gender Discrimination in Transportation Network Companies," *NBER Working Paper Series*, 1–38.

Glover, D., A. Pallais, and W. Pariente (2017): "Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores," *The Quarterly Journal of Economics*, 132, 1219–1260.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997): "Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *The Review of Economic Studies*.

Hjort, J. (2014): "Ethnic divisions and production in firms," *The Quarterly Journal of Economics*, 129, 1899–1946.

Holmström, B. (1999): "Managerial incentive problems: A dynamic perspective," *The Review of Economic Studies*, 66, 169–182.

Iacus, S. M., G. King, and G. Porro (2009): "CEM : Software for Coarsened Exact Matching," *Journal of Statistical Software*.

Kakar, V., J. Voelz, J. Wu, and J. Franco (2018): "The visible host: Does race guide Airbnb rental rates in San Francisco?" *Journal of Housing Economics*, 40, 25–40.

KUZNETS, S. (1955): "Economic growth and income inequality," *The American economic review*, 45, 1–28.

LAMBIN, X. AND E. PALIKOT (2020): "The Impact of Online Reputation on Ethnic Discrimination," *TSE Mimeo*.

LAOUENAN, M. AND R. RATHELOT (2017): "Ethnic Discrimination on an Online Marketplace of Vacation Rentals," *working paper*.

MAZZELLA, F. AND A. SUNDARARAJAN (2016): "Entering the Trust Age," Tech. rep.

MCFADDEN, D. (1974): "Conditional logit analysis of qualitative choice behavior," *Frontiers in Econometrics*.

PALLAIS, A. (2014): "Inefficient hiring in entry-level labor markets," *American Economic Review*, 104, 3565–99.

PHELPS, E. S. (1972): "The Statistical theory of Racism and Sexism," *American Economic Review*, 62, 659–661.

REYNAERT, M. AND F. VERBOVEN (2014): "Improving the performance of random coefficients demand models: the role of optimal instruments," *Journal of Econometrics*, 179, 83–98.

ROSENBAUM, P. R. AND D. B. RUBIN (1983): "The central role of the propensity score in observational studies for causal effects," *Biometrika*.

SARSONS, H. (2017): "Interpreting Signals in the Labor Market: Evidence from Medical Referrals," *Job Market Paper*.

TADELIS, S. (2016): "Reputation and feedback systems in online platform markets," *Annual Review of Economics*, 8, 321–340.

TJADEN, J. D., C. SCHWEMMER, AND M. KHADJAVI (2018): "Ride with Me -Ethnic Discrimination, Social Markets, and the Sharing Economy," *European Sociological Review*, 34, 418–432.

ZERVAS, G., D. PROSERPIO, AND J. BYERS (2015): “A First Look at Online Reputation on Airbnb, Where Every Stay is Above Average,” *Working paper*, 1–22.

2.9 Appendix

Navigation on Blablacar.fr

First, users type in the origin, destination and date of the ride they are seeking. They then see a list of rides meeting their request (Figure 2.14). They may then click on specific postings to have more details about the ride (Figure 2.15). Finally they may either see the profile of the driver (Figure 2.16) or proceed directly to payment. BlaBlaCar service fees are a function of the price posted by the driver. The fees and their evolution over time are shown on Figure 2.17.

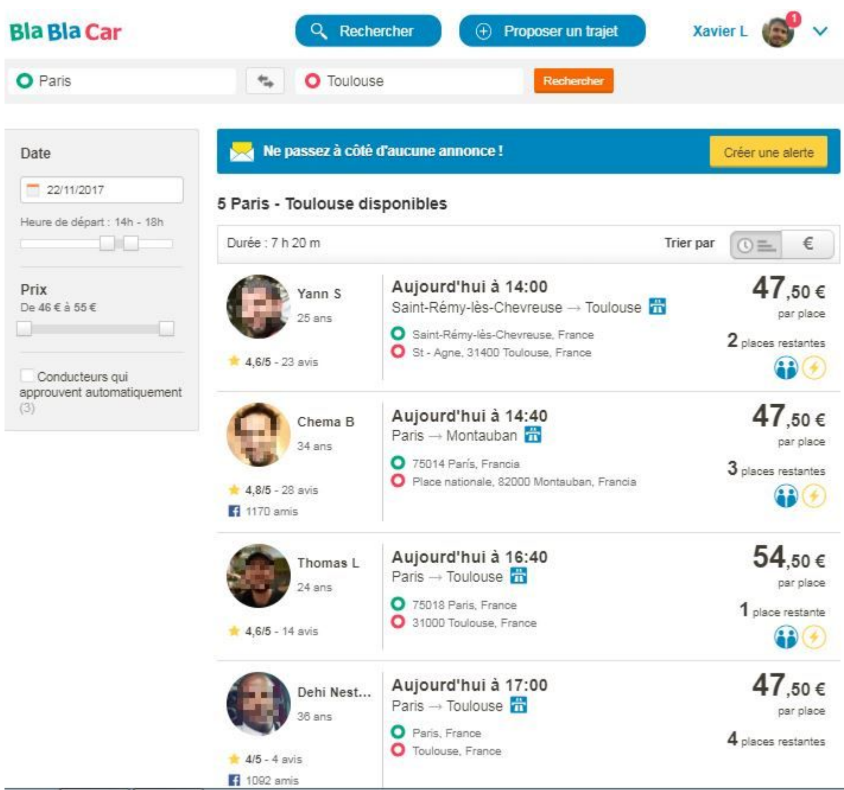

Figure 2.14: Listing offered on a given route

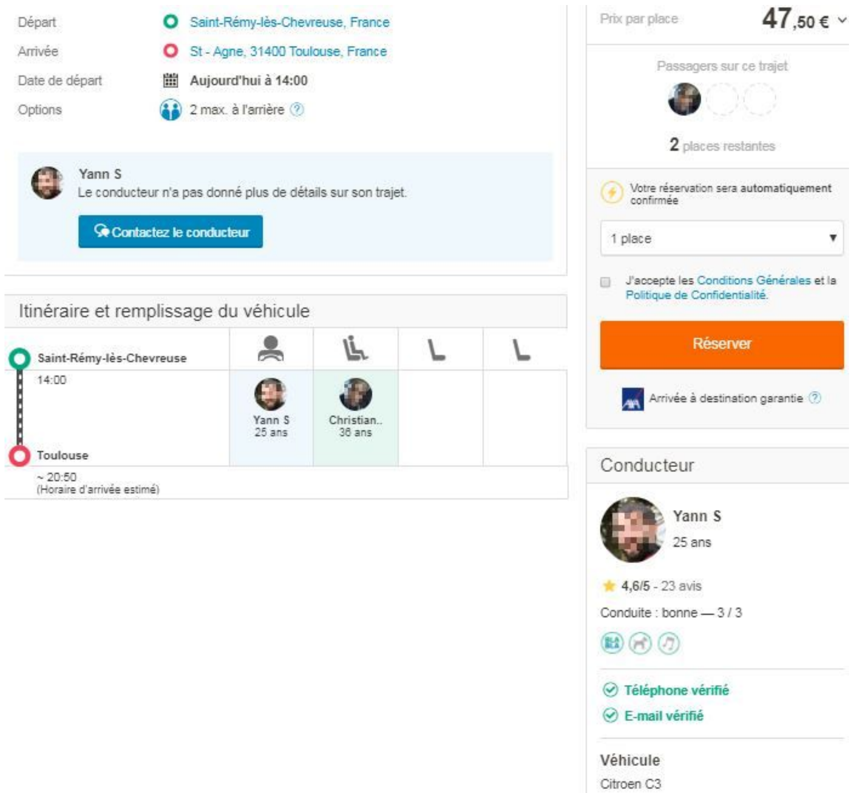

Figure 2.15: Details of a posting

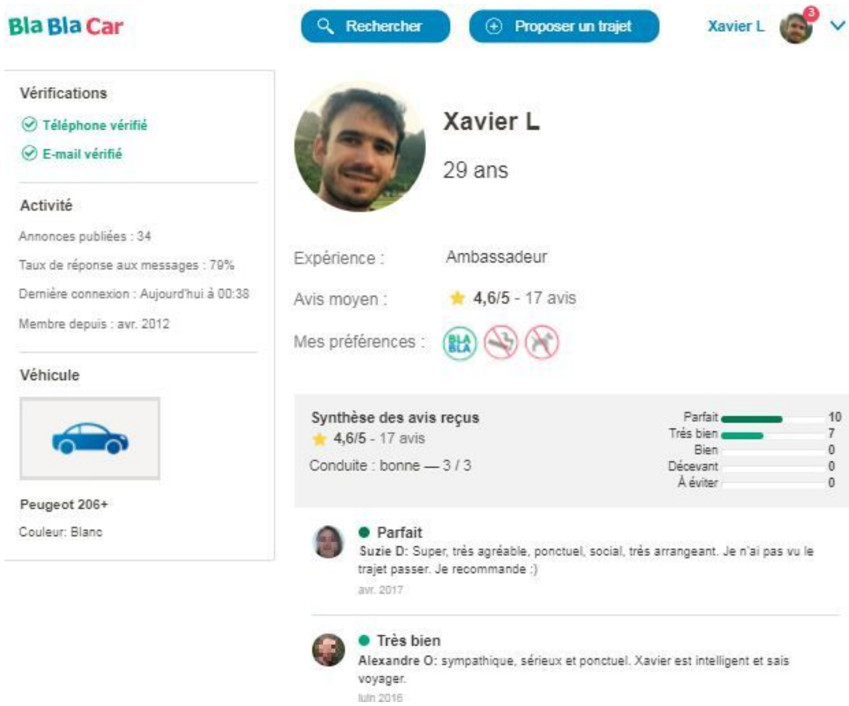

Figure 2.16: A driver's profile

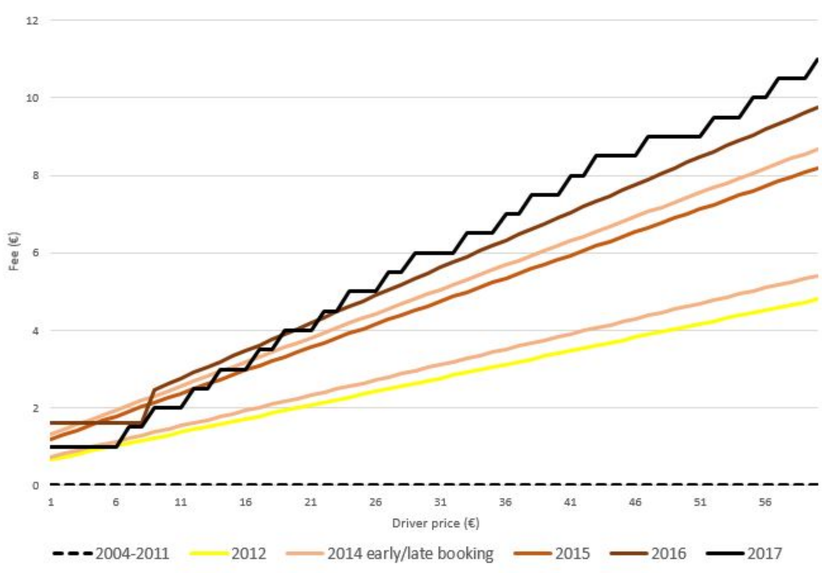

Figure 2.17: Evolution of service fees on BlaBlaCar over time

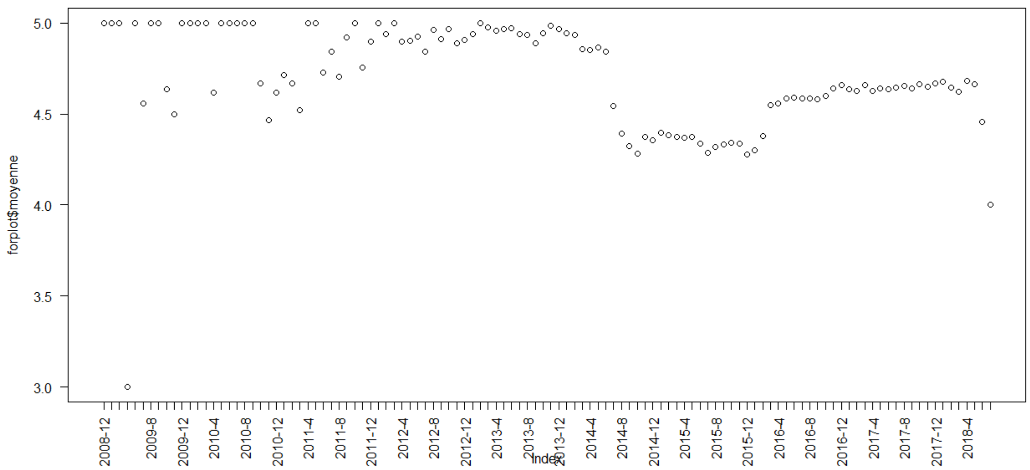

Figure 2.18: Average rating for drivers with more than 30 reviews

Changes in the BlaBlaCar reputation system

In our study of the evolution of ratings, we have abstracted from the potential changes in the design of the reputation system of BlaBlaCar. Some drivers in our sample have been BlaBlaCar users since December 2008, and others joined only a few days before our crawler observes their listing. These drivers may have operated under different market characteristics. See Figure 2.18 for the evolution of the average rating over time. Until the end of 2013, ratings were either 1 or 5. In early 2014, these binary ratings were translated to the current 5-star system. Later, in February 2016, the wording of the ratings was changed: *excellent* became *tres bien* and *extraordinaire* became *parfait*. The impact of this change on the average rating is clear. People are more likely to call a ride *parfait* than they were to call it *extraordinaire*. Finally, these changes influenced the informativeness of the reputation system; see Figure 2.19. The dotted black line shows HHI (which is a measure of dispersion and, hence, the informativeness of the classifiers): the smaller the HHI is, the more informative the classifier. The ratings in the period 2014-2016 were the most informative. Dark green, green, orange, pink, and red represent the shares of 5s, 4s, 3s, 2s and 1s, respectively. Initially, there is a considerable noise

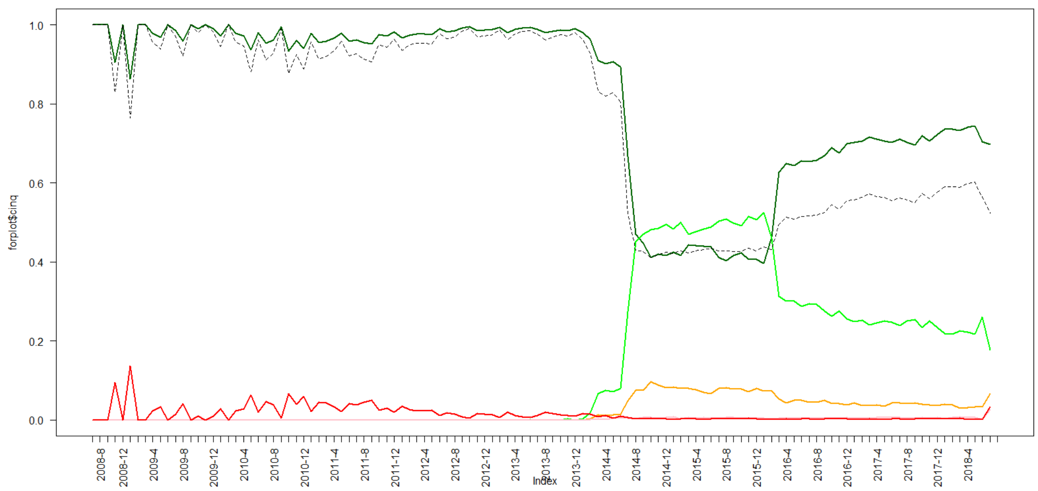

Figure 2.19: Informativeness of the reputation system and share of grades received. Dark green=5, light green=4, orange=3, pink=2, red=1.

because we have very few observations: fewer than 100 per month before October 2009 and more than 30.000 per month starting in 2017.
These changes are important because they affected the ratings that we study, but they also show how important the design of the review system is. One may be concerned that some of the decline in effort that we characterized could be due to changes in the reputation system. In a sample restricted to drivers who joined after all the changes in the reputation system were made, we can reproduce the same patterns of behavior; however, we lose a considerable number of observations. Thus, we argue that the evolution of ratings throughout the career of a driver on BlaBlaCar is due to the economic logic of career concerns rather than exogenous changes in the reputation system.

Classification method for gender and ethnicity

Driver-specific characteristics are key determinants in our model. Hence, the drivers' type must be identified as accurately as possible. Specifically, gender and ethnicity are critical to our analysis. To identify these characteristics, both prospective riders and the econometrician consider two relevant sources of information: the first name and the profile picture. We use both information to infer gender and ethnicity.

Classification of gender

As a first source of information, we use the name of the driver. We match our dataset of driver names with those of various sources relating first names with ethnicity. The French Government repository of names (www.data.gouv.fr/fr/datasets/liste-de-prenoms) constitutes our main source of information. We complement it with data from other sources.[39] This data enables us to identify the gender of almost 80% of drivers, along with 3% unisex names.

We then use facial recognition to identify gender whenever a picture is available. This process also enable us to identify 80 % of the dataset. By combining these two processes, we can directly identify gender for 95% of the dataset.

Further, we use facial recognition to enrich and correct our name database. Rare or misspelled names (either because the driver registered under a nick-name or because of translation variations if the name is not originally French) can be re-classified. This process can identify the gender of some drivers whose names are not listed in our inventories and who do not have a picture (or for pictures where gender is not easily identified) because other drivers with the same name may have posted identifiable pictures. This method brings the precision of our gender identification as high as 99%. Figure 2.20 summarizes our identification process.

[39] www.signification-prenom.net, www.madame.lefigaro.fr/prenoms/origine

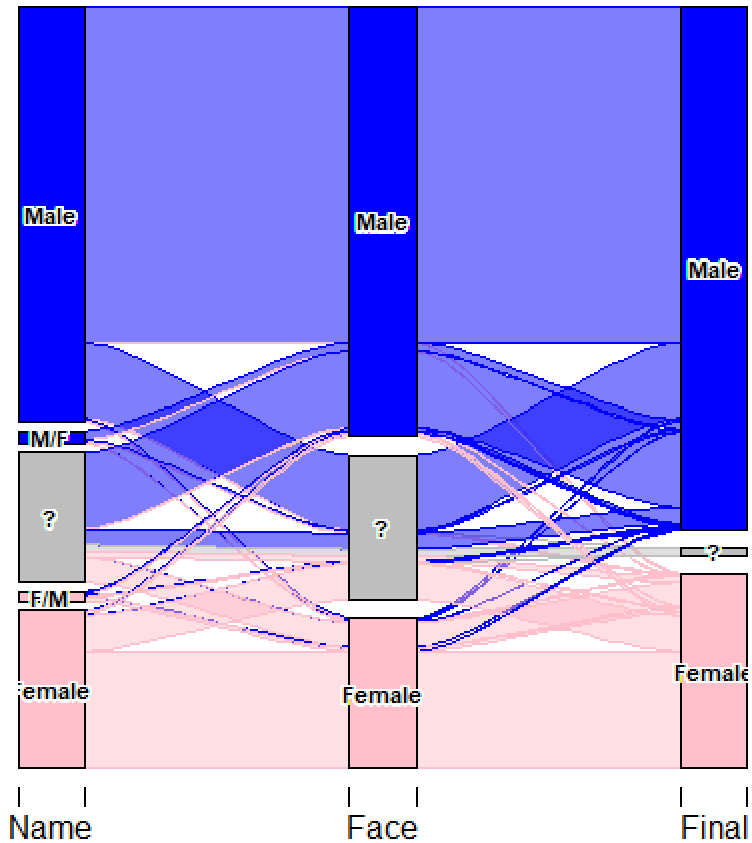

Figure 2.20: Classification process for gender: by name (left), by facial recognition (center) and final classification (right)

Classification of ethnicity

Our methodology for the identification of ethnicity follows the same steps and uses the same sources as those for gender classification. First, we collect the origins of names from the data sources mentioned above. This provides the ethnicity of approximately 81% of our sample. However, names might not be a perfect indicator of ethnicity. Indeed, many visible minorities have a French name for various historical reasons or because they have foreign origins but were born in France. In that case, a simple name analysis would classify them as non-minorities while they might belong to a minority on the basis of their skin color.

Hence, we use facial recognition to identify ethnicity whenever a picture is available. The algorithm proposes an ethnicity for 80 % of the dataset. However, only “white”, “black”, “Asian” , and “Latino” ethnicities are proposed. People of Arabic origin are classified as “white”. Hence, facial recognition is useful only to classify drivers more accurately between african origin, and majority or arabic origin.

We also use facial recognition to enrich and correct our name repository and to better identify ethnicity. Overall, facial recognition reclassifies 2.5% of drivers with a French name and 5% of drivers with Arabic names (predominantly Muslim names) into Sub-Saharan ethnicity. Including facial recognition increases the sample size for minorities from 11% to 14% of our sample. Figure 2.21 summarizes our identification process.

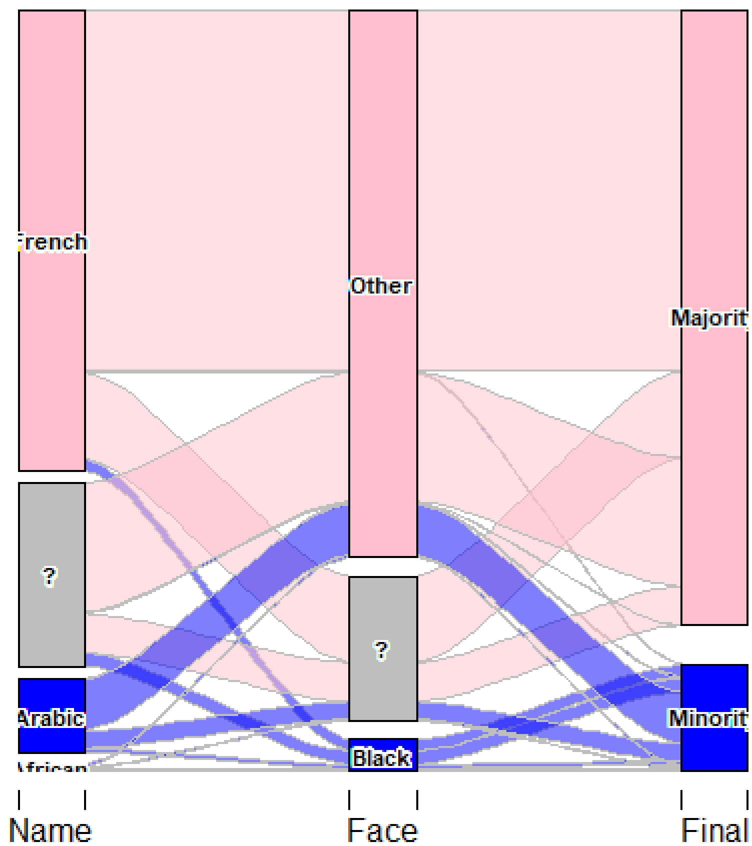

Figure 2.21: Classification process for ethnicity: by name analysis (left), by facial recognition (center) and final classification (right)

Ratings as a measure of passenger satisfaction

The body of the paper analyses the effect of reputation on the sole basis of ratings. It assumes that ratings have enough informational content to allow passengers to form a belief about the quality of a driver.

In this Appendix, we show that ratings are indeed likely to be a good summary of passengers' experience. To do so, we analyze whether good reviews (i.e. reviews with a high rating) are more likely to be associated with a written comment that has a positive connotation than bad reviews. For that purpose, we use the Cloud Natural Language processing tools of Google, a tool that uses machine learning to reveal the structure and meaning of text. We are particularly interested in the sentiment of the review, with a measure between -1 (very negative) and 1 (very positive).

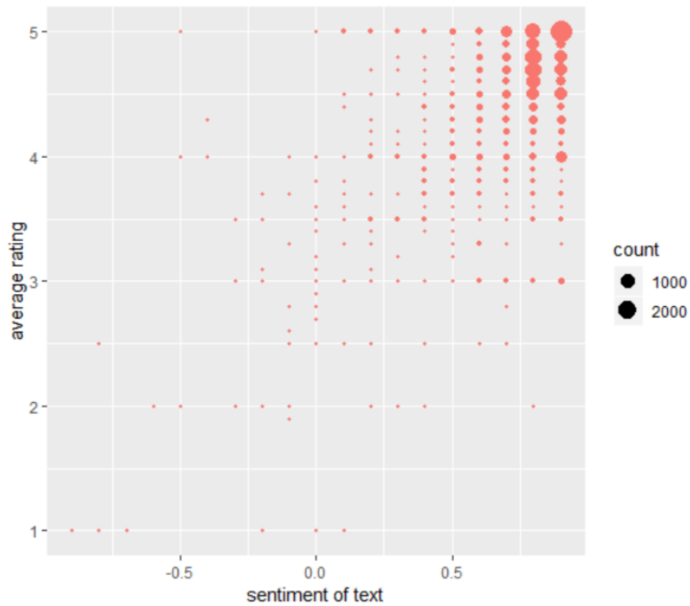

Figure 2.22: Textual analysis of the comments (18000 randomly selected drivers). Average rating and average sentiment of written comment are highly correlated.

The correlation between the grade given, and the sentiment of the text of the review very high, as is suggested by Figure 2.22. We therefore conclude that ratings are a satisfactory a measure of performance for the purpose of the present paper.

Oversampling of minorities for short-notice rides

Due to our scraping method, it cannot be excluded that our sample provides a slightly biased representation of listings. Indeed, the program takes snapshots of listings displayed on the website at a given point time. However, rides that are already full are no longer displayed on the platform. This means our data collection may undersample the particularly attractive rides that would sell out very fast, or those corresponding to times when demand is much higher than supply. This wouldn't be an issue if both minorities and non-minorities were affected the same way by this sampling bias. However, as we show in this paper the minority status does impact the attractiveness of a given listing. Therefore, minorities who may be perceived as posting less attractive rides remain longer on display and may therefore be over-represented in our sample. Therefore, our minority gap estimates should be understood as lower bounds. Indeed, minorities are compared to a pool constituted of non-minorities that are not so good as to have sold out their seats extremely fast. Table 2.23 shows that minority drivers represent a specially high share of rides that are posted on a short notice, a possible sign that non-minority drivers have sold their seats faster. For trips posted with more notice, we believe our sample is indeed representative of the actual participants on blablacar. Indeed, most of the rides –either from minorities or not – still have more than one empty seat, which means that most listings and indeed collected. In fact, Blablacar informs drivers that most passengers book rides only a few days in advance.

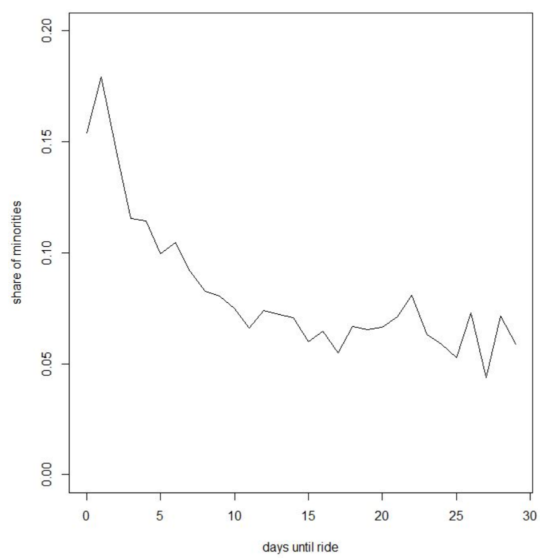

Figure 2.23: Share of minorities in sample as a function of number of days between posting and departure

This is true despite the fact minorities tend to allow for automatic confir-

name of a variable	description
price	price set by the driver in EUR; has to be lower than maximum price: 0.082 per km
age	age of the driver in years
reviews	number of reviews received by the driver
male	gender defined based on photo recognition and name
minority	takes the value of one when the driver is of Arabic or African origin, and zero otherwise; defined based on photo recognition and name (see. Lambin& Palikot (2019)) for details)
picture	takes the value of one when driver added a picture, and zero otherwise
talkative	categorical variable (bla, blabla, blablabla) indicating how talkative the driver is
bio	number of words in driver's description
ride description	number of words in ride's description
reputation	mean of grades received by the driver
published rides	number of rides ever published by the driver
number of clicks	number of clicks a given listing has received; clicking is necessary for booking a ride but not sufficient; measured at the moment of data collection
sold seats	number of seats already sold; measured at the moment of data collection
revenue	sold seats multiplied by price
posts per month	mean number of listings posted by the driver since she joined the platform
seniority	number of months since the driver joined the platform
competition	number of listings available on the same day on the same route
median revenue	mean of median revenues in cities of departure and arrival; source: INSEE
public transport	travelling time by public transport on the route at listings' departure time; source: Google API
train strike	SNCF official strike implicating a given route
value of car	price of a comparable car model in thousands of EUR; when a model of a car is not available mean price of a brand; source: ebay (scrapped data)
fuel consumption	mean fuel consumption of a model of a car; when model of a car is not available mean consumption of a brand; source: ADEME
length (km)	distance in km between cities of departure and arrival; souce: Google API
lengh (hours)	estimated driving time by a car on a given route and time; source: Google API
hours until departure	number of hours between data collection and a ride departure
posted since	number of hours between the posting of the listing and data collection
automatic acceptance	takes the value of one if booking requests are automatically accepted and zero if the driver chose to accept/reject requests manually
to fuel price	average price of a litre of diesel in a city of arrival in cents
from fuel price	average price of a litre of diesel in a city of departure in cents
toll viamich	total toll costs on a given route in EUR; source: https://www.viamichelin.com/
travel costs	mean of fuel costs multiplied by fuel consumption plus toll fees
weekday	takes a value of 1 on weekdays and zero on weekends
pets	takes a value of 1 if the driver accepts pets and zero otherwise
music	takes a value of 1 if the driver listens to music in the car and zero otherwise
smoke	takes a value of 1 if the driver accepts smoking in the car and zero otherwise
detour	categorical variable: 1 if no detour, 2 if some detour (up to 15 min), and 3 if more than 15 minutes detour
luggage	categorical variable: 1 if no luggage, 2 if small bags, 3 if big bags are allowes

Table 2.9: Defintion of main variables

mation more frequently than non-minorities (18% of drivers with automatic confirmation are minorities, while they represent only 12% of the drivers with manual confirmation).

Definition of variables

Sources of supplementary data

- Databaset of names constructed based on: French government statistics www.data.gouv.fr/fr/datasets/liste-de-prenoms and supplemented with www.signification-prenom.net , www.madame.lefigaro.fr/prenoms/origine

- Car prices on eBay Germany: www.kaggle.com/orgesleka/used-cars-database
- Fuel consumtpion of cars: French environment and energy management agency- ADEME
- City specific population, median income, index of crime, and a share of foreign born residents- French statistics office INSEE.

Output gap: endogeneity of price

In this section, we address the problem of endogeneity of price and quantity in the regression showing minority output gap. Column 1 of the Table 2.10 introduced the price of the ride in the regression with sold seats as a left-hand side variable (other covariates are unchanged). Column 2 presents an instrumental variables regression, where the price is instrumented with a price of car fuel in the cities of departure and arrival (which we observe on the daily basis), and highway tolls on a given route in a given period.

Reputation effect

Panel data results

Thousands of drivers are active on BlaBlaCar at any moment; thus, every time we collect data, we observe only a fraction of all available listings. As a consequence, we see most drivers only once. However, in some cases (22.800 drivers), we see the driver at least twice, which gives us a panel with almost 56.800 observations. However, this sample is unbalanced, with drivers being observed between 2 and 30 times. We use several standard models that allow us to compare the gap associated with being a minority entrant or incumbent entrant. Reduction in the sample size results in lower significance of our estimates. However, the signs and point estimates appear to confirm our hypothesis.

Table 2.10: Sold seats: controlling for price and instrumenting it.

	Dependent variable:	
	sold seats	
	OLS	*IV*
	(1)	(2)
minority	−0.013*** (0.003)	−0.006** (0.003)
price	−0.009*** (0.0002)	−0.024*** (0.002)
driver age	−0.001*** (0.0001)	−0.001*** (0.0001)
reviews	0.001*** (0.00005)	0.001*** (0.0001)
reviews2	−0.00000*** (0.00000)	−0.00000*** (0.00000)
male	−0.0002 (0.002)	−0.005** (0.002)
hours untill ride	−0.001*** (0.00001)	−0.001*** (0.00001)
posted since	0.012*** (0.0002)	0.012*** (0.0002)
post per month	−0.005*** (0.001)	−0.005*** (0.001)
length bio	0.0002 (0.0001)	−0.0002* (0.0001)
car price	0.0002 (0.0002)	0.001** (0.0002)
public transport ratio	10.022*** (3.365)	17.286*** (3.499)
km	0.001*** (0.00005)	0.002*** (0.0001)
day	0.018*** (0.004)	0.015*** (0.004)
night	−0.054*** (0.006)	−0.050*** (0.007)
train strike	0.131*** (0.007)	0.171*** (0.006)
length ride (# words)	0.001*** (0.0001)	0.0005*** (0.0001)
picture	0.001 (0.006)	0.002 (0.006)
automatic acceptance	0.114*** (0.002)	0.086*** (0.004)
weekday	−0.042*** (0.004)	−0.041*** (0.004)
day*weekday	0.009* (0.005)	0.012** (0.005)
night*weekday	0.004 (0.008)	0.004 (0.008)
Constant	0.202*** (0.041)	0.046 (0.043)
Observations	318,420	287,754
R^2	0.078	0.064
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 2.11: Revenue regressed over driver and ride characteristics

	Dependent variable:		
	revenue		
	(1)	(2)	(3)
minority	-0.623^{***} (0.142)	-0.451^{**} (0.178)	-0.233 (0.168)
driver age	-0.033^{***} (0.004)	-0.026^{***} (0.005)	-0.008 (0.005)
reviews	0.234^{***} (0.029)	0.078^{***} (0.021)	0.014^{***} (0.001)
male	-0.278^{**} (0.108)	-0.198 (0.131)	0.213 (0.151)
hours till ride	-0.015^{***} (0.0005)	-0.019^{***} (0.001)	-0.028^{***} (0.001)
posted since	0.215^{***} (0.007)	0.304^{***} (0.009)	0.345^{***} (0.009)
post per month	0.043 (0.040)	-0.133^{***} (0.049)	-0.409^{***} (0.034)
seniority (# months)	-0.006^{***} (0.002)	-0.016^{***} (0.003)	-0.029^{***} (0.003)
length bio	0.009 (0.005)	0.010 (0.006)	-0.00002 (0.007)
car price	-0.018^{*} (0.010)	-0.006 (0.012)	-0.019 (0.012)
competition	0.006^{***} (0.002)	0.004^{*} (0.002)	0.005^{***} (0.002)
duration public transport	-0.192 (0.534)	-0.930 (0.678)	-2.719^{***} (0.804)
km	0.003 (0.003)	0.010^{***} (0.004)	0.017^{***} (0.004)
day	0.402^{**} (0.194)	0.770^{***} (0.239)	0.574^{**} (0.246)
night	-1.025^{***} (0.299)	-0.860^{**} (0.391)	-1.668^{***} (0.378)
train strike	2.757^{***} (0.294)	2.981^{***} (0.358)	3.107^{***} (0.547)
length ride (# words)	0.029^{***} (0.004)	0.018^{***} (0.004)	0.013^{***} (0.004)
picture	0.199 (0.278)	0.091 (0.432)	-1.296^{***} (0.413)
automatic acceptance	3.381^{***} (0.107)	3.126^{***} (0.126)	2.929^{***} (0.126)
weekday	-0.509^{**} (0.202)	-0.222 (0.248)	-1.173^{***} (0.244)
travel cost	0.017^{**} (0.008)	0.005 (0.011)	0.020 (0.012)
median revenue	-0.032 (0.130)	-0.061 (0.161)	-0.184 (0.184)
day*weekday	0.096 (0.238)	-0.237 (0.291)	0.243 (0.293)
night*weekday	0.013 (0.368)	-0.612 (0.474)	0.359 (0.454)
Constant	3.745^{**} (1.894)	4.438^{*} (2.466)	15.360^{***} (3.153)
Driver effects	X	X	X
Ride effects	X	X	X
Time effects	X	X	X
Trip effects	X	X	X
Observations	82,563	65,013	68,505
R^2	0.060	0.070	0.096
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2.12: Sold seats regressed over driver and ride characteristics

	Dependent variable:		
	taken_seats		
	(1)	(2)	(3)
minority	−0.024*** (0.005)	−0.014** (0.006)	−0.0003 (0.007)
driver age	−0.001*** (0.0001)	−0.001*** (0.0002)	−0.0004** (0.0002)
reviews (#)	0.008*** (0.001)	0.004*** (0.001)	0.001*** (0.00004)
male	−0.005 (0.004)	−0.002 (0.004)	0.011* (0.006)
seniority (# months)	−0.0002** (0.0001)	−0.001*** (0.0001)	−0.001*** (0.0001)
hours till ride	−0.001*** (0.00002)	−0.001*** (0.00002)	−0.001*** (0.00003)
posted since	0.008*** (0.0002)	0.012*** (0.0003)	0.014*** (0.0004)
post per month	0.001 (0.001)	−0.007*** (0.002)	−0.020*** (0.001)
length bio	0.0001 (0.0002)	0.0001 (0.0002)	−0.0001 (0.0003)
car price	−0.0002 (0.0003)	0.0002 (0.0004)	−0.00003 (0.0004)
competition	0.0002*** (0.0001)	0.0001** (0.0001)	0.0002** (0.0001)
public transport ratio	10.958** (5.263)	1.612 (6.970)	10.991 (9.310)
km	−0.00004 (0.0001)	−0.0001 (0.0001)	0.00003 (0.0001)
day	0.013** (0.006)	0.024*** (0.008)	0.019** (0.010)
night	−0.044*** (0.010)	−0.030** (0.013)	−0.079*** (0.015)
train strike	0.103*** (0.010)	0.139*** (0.013)	0.151*** (0.022)
length ride(# words)	0.001*** (0.0001)	0.001*** (0.0001)	0.001*** (0.0001)
picture	−0.002 (0.009)	−0.007 (0.014)	−0.026* (0.015)
automatic acceptance	0.134*** (0.004)	0.125*** (0.004)	0.134*** (0.005)
weekday	−0.029*** (0.007)	−0.028*** (0.009)	−0.074*** (0.009)
day*weekday	0.010 (0.008)	−0.002 (0.010)	0.025** (0.011)
night*weekday	0.005 (0.012)	−0.024 (0.016)	0.022 (0.018)
Constant	0.137** (0.060)	0.279*** (0.084)	0.470*** (0.124)
Observations	91,870	72,597	76,999
R^2	0.066	0.066	0.083
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2.13: Number of clicks regressed over driver and ride characteristics

	Dependent variable:		
	number of clicks		
	(1)	(2)	(3)
minority	−0.472*** (0.155)	−0.376** (0.176)	0.012 (0.157)
driver age	−0.074*** (0.004)	−0.062*** (0.005)	−0.033*** (0.005)
reviews (#)	−0.025 (0.031)	0.052** (0.020)	0.017*** (0.001)
male	−1.490*** (0.118)	−1.741*** (0.129)	−0.917*** (0.140)
seniority (# months)	−0.007*** (0.002)	−0.016*** (0.003)	−0.044*** (0.003)
hours till ride	−0.034*** (0.001)	−0.038*** (0.001)	−0.045*** (0.001)
posted since	1.195*** (0.008)	1.348*** (0.010)	1.210*** (0.009)
post per month	−0.186*** (0.044)	−0.470*** (0.048)	−0.780*** (0.032)
length bio	0.002 (0.006)	0.005 (0.006)	−0.004 (0.006)
car price	0.013 (0.010)	0.016 (0.012)	−0.0004 (0.011)
competition	0.011*** (0.002)	0.009*** (0.002)	0.012*** (0.002)
public transport ratio	621.914*** (172.633)	462.974** (201.728)	210.319 (227.733)
km	0.016*** (0.002)	0.016*** (0.002)	0.014*** (0.003)
day	−0.711*** (0.209)	0.494** (0.236)	0.629*** (0.227)
night	−0.078 (0.328)	0.593 (0.386)	−1.286*** (0.351)
train strike	5.235*** (0.330)	4.815*** (0.364)	4.839*** (0.533)
length ride (# words)	0.052*** (0.004)	0.037*** (0.004)	0.011*** (0.003)
picture	1.228*** (0.282)	0.236 (0.389)	−0.883** (0.355)
automatic acceptance	−0.194* (0.116)	−0.454*** (0.124)	−1.561*** (0.118)
weekday	−1.253*** (0.218)	−0.146 (0.245)	−0.894*** (0.225)
day*weekday	1.395*** (0.257)	0.059 (0.287)	0.664** (0.271)
night*weekday	−0.160 (0.403)	−0.591 (0.469)	0.846** (0.423)
Constant	6.853*** (1.990)	7.978*** (2.446)	18.422*** (3.060)
Observations	87,004	69,163	73,834
R^2	0.250	0.259	0.254
Note:			*p<0.1; **p<0.05; ***p<0.01

	Dependent variable: number of clicks		
	Pooled	Between	Random
minority	0.288 (0.202)	0.409 (0.275)	0.317 (0.236)
entrant	-0.995^{***} (0.143)	-0.811^{***} (0.179)	-0.764^{***} (0.155)
minority*entrant	-0.678^{*} (0.353)	-0.692 (0.449)	-0.717^{*} (0.387)
driver's age	-0.036^{***} (0.005)	-0.038^{***} (0.006)	-0.036^{***} (0.006)
talkative	0.220^{*} (0.123)	0.363^{**} (0.156)	0.282^{**} (0.141)
male	-1.074^{***} (0.142)	-1.105^{***} (0.171)	-1.128^{***} (0.159)
hours until ride	-0.028^{***} (0.0005)	-0.023^{***} (0.001)	-0.029^{***} (0.0005)
posted since	1.136^{***} (0.010)	1.068^{***} (0.016)	1.172^{***} (0.010)
bio (# words)	-0.002 (0.004)	-0.003 (0.005)	-0.002 (0.004)
car price	-0.018 (0.012)	-0.031^{**} (0.015)	-0.021 (0.014)
competition	0.036^{***} (0.002)	0.035^{***} (0.003)	0.034^{***} (0.002)
median revenue	-0.00002 (0.00003)	-0.0001^{*} (0.00004)	-0.00000 (0.00003)
public transport ratio	-0.909 (7.222)	-1.761 (10.861)	-2.131 (7.765)
km	0.007^{***} (0.0004)	0.006^{***} (0.001)	0.006^{***} (0.0004)
day	0.538^{**} (0.231)	0.574 (0.364)	0.462^{**} (0.231)
night	-0.605^{*} (0.357)	-1.134^{*} (0.581)	-0.763^{**} (0.358)
train strike	3.269^{***} (0.325)	3.049^{***} (0.538)	3.545^{***} (0.319)
ride (# words)	0.018^{***} (0.002)	0.021^{***} (0.002)	0.020^{***} (0.002)
picture	0.246 (0.201)	0.494^{*} (0.260)	0.496^{**} (0.230)
automatic acceptance	-1.334^{***} (0.122)	-1.299^{***} (0.164)	-1.307^{***} (0.132)
weekday	-0.018 (0.236)	-0.457 (0.387)	0.112 (0.237)
consumption	0.278^{***} (0.084)	0.377^{***} (0.106)	0.303^{***} (0.095)
day*weekday	0.465 (0.284)	0.925^{**} (0.460)	0.389 (0.284)
night*weekday	-0.018 (0.444)	1.577^{**} (0.746)	0.094 (0.443)
Constant	11.748^{***} (0.880)	10.465^{***} (1.234)	11.477^{***} (0.948)
Observations	56,760	22,794	56,760
R^2	0.244	0.220	0.262
Note:			*p<0.1; **p<0.05; ***p<0.01

We estimate the following model:

$$y_{it} = \alpha + X_{it}\beta + Z_i\gamma + c_i + \tau_t + \epsilon_{it}$$

where c_i are individual fixed effects and ϵ_{it} is an idiosyncratic error term.

We present minority dummies and the products of minority and entrant dummies. Similarly to the cross-sectional analysis in the main body of the paper, we conclude that upon entering the market, minority drivers receive lower outcomes and that this effect weakens as drivers receive reviews. Again, the reputation effect is significant for all measures of economic performance.

Strikes

Table 2.14 presents means of selected characteristics of drivers on days of strike and days without a strike. Subset of drivers active in the period 03/04/2018 to 28/06/2018.

Results of the main specification with number of sold seats as the dependent

	Dependent variable: sold seats		
	Pooled	Between	Random
minority	0.002 (0.009)	0.016 (0.011)	0.002 (0.009)
entrant	-0.060^{***} (0.011)	-0.058^{***} (0.012)	-0.059^{***} (0.011)
minority*entrant	-0.035^{**} (0.016)	-0.041^{**} (0.018)	-0.035^{**} (0.016)
male	0.005 (0.008)	0.004 (0.009)	0.004 (0.008)
driver's age	-0.0004^{*} (0.0002)	-0.0003 (0.0002)	-0.0004^{*} (0.0002)
talkative	0.001 (0.005)	0.003 (0.006)	0.001 (0.006)
hours until ride	-0.001^{***} (0.00002)	-0.001^{***} (0.00003)	-0.001^{***} (0.00002)
posted since	0.016^{***} (0.0004)	0.012^{***} (0.001)	0.016^{***} (0.0004)
bio (# words)	-0.0001 (0.0002)	-0.0003 (0.0002)	-0.0001 (0.0002)
car price	-0.0003 (0.001)	-0.001 (0.001)	-0.0004 (0.001)
competition	0.001^{***} (0.0001)	0.001^{***} (0.0001)	0.001^{***} (0.0001)
median revenue	0.00000^{***} (0.00000)	0.00001^{***} (0.00000)	0.00000^{***} (0.00000)
public transport ratio	-0.146 (0.318)	-0.566 (0.440)	-0.147 (0.322)
km	-0.00002 (0.00002)	-0.0001^{**} (0.00002)	-0.00002 (0.00002)
day	0.015 (0.010)	0.004 (0.015)	0.015 (0.010)
night	-0.048^{***} (0.016)	-0.062^{***} (0.023)	-0.048^{***} (0.016)
train strike	0.126^{***} (0.014)	0.110^{***} (0.022)	0.128^{***} (0.014)
ride (# words)	0.0004^{***} (0.0001)	0.0004^{***} (0.0001)	0.0004^{***} (0.0001)
picture	0.002 (0.009)	0.015 (0.010)	0.003 (0.009)
automatic acceptance	0.109^{***} (0.005)	0.108^{***} (0.007)	0.109^{***} (0.005)
weekday	-0.045^{***} (0.010)	-0.059^{***} (0.016)	-0.045^{***} (0.010)
consumption	0.020^{***} (0.004)	0.021^{***} (0.004)	0.020^{***} (0.004)
minority*entrant	-0.035^{**} (0.016)	-0.041^{**} (0.018)	-0.035^{**} (0.016)
entrant*male	-0.016 (0.013)	-0.022 (0.014)	-0.015 (0.013)
day*weekday	0.019 (0.012)	0.034^{*} (0.019)	0.019 (0.013)
night*weekday	-0.020 (0.019)	0.012 (0.030)	-0.020 (0.020)
Constant	0.180^{***} (0.039)	0.158^{***} (0.050)	0.175^{***} (0.039)
Observations	59,359	23,076	59,359
R^2	0.089	0.085	0.088
Note:			*p<0.1; **p<0.05; ***p<0.01

	Dependent variable: revenue		
	Pooled	Between	Random
minority	-0.334 (0.213)	0.022 (0.275)	-0.272 (0.228)
entrant	-1.387^{***} (0.150)	-1.452^{***} (0.179)	-1.308^{***} (0.155)
minority*entrant	-0.680^{*} (0.372)	-0.866^{*} (0.448)	-0.741^{*} (0.387)
driver's age	-0.006 (0.005)	-0.002 (0.006)	-0.005 (0.005)
talkative	0.020 (0.129)	0.065 (0.155)	0.026 (0.137)
male	-0.201 (0.148)	-0.307^{*} (0.170)	-0.240 (0.156)
hours untill ride	-0.018^{***} (0.0005)	-0.016^{***} (0.001)	-0.019^{***} (0.0005)
posted since	0.371^{***} (0.010)	0.290^{***} (0.014)	0.375^{***} (0.010)
bio (# words)	-0.001 (0.004)	-0.005 (0.005)	-0.001 (0.004)
car price	-0.007 (0.013)	-0.022 (0.015)	-0.010 (0.013)
competition	0.024^{***} (0.002)	0.024^{***} (0.003)	0.024^{***} (0.002)
median revenue	0.0002^{***} (0.00003)	0.0003^{***} (0.00004)	0.0002^{***} (0.00003)
public transport ratio	-33.375^{***} (7.569)	-40.181^{***} (10.934)	-33.318^{***} (7.835)
km	0.013^{***} (0.0004)	0.011^{***} (0.001)	0.013^{***} (0.0004)
day	0.410^{*} (0.243)	0.445 (0.367)	0.400 (0.244)
night	-1.341^{***} (0.373)	-2.091^{***} (0.579)	-1.300^{***} (0.376)
train strike	2.367^{***} (0.339)	1.779^{***} (0.543)	2.429^{***} (0.338)
ride (# words)	0.007^{***} (0.002)	0.008^{***} (0.002)	0.007^{***} (0.002)
picture	0.087 (0.212)	0.366 (0.260)	0.170 (0.225)
automatic acceptance	2.064^{***} (0.128)	2.012^{***} (0.164)	2.104^{***} (0.133)
weekday	-0.847^{***} (0.249)	-1.142^{***} (0.390)	-0.828^{***} (0.251)
consumption	0.315^{***} (0.088)	0.340^{***} (0.106)	0.325^{***} (0.093)
minority*entrant	-0.680^{*} (0.372)	-0.866^{*} (0.448)	-0.741^{*} (0.387)
day*weekday	0.317 (0.299)	0.446 (0.465)	0.290 (0.300)
night*weekday	-0.215 (0.465)	0.938 (0.745)	-0.232 (0.467)
Constant	-1.089 (0.926)	-2.287^{*} (1.239)	-1.200 (0.957)
Observations	58,621	23,018	58,621
R^2	0.095	0.093	0.094
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2.14: Characteristics of drivers on days of strike and non-strike days.

	Minority	Male	Reviews	Notice	Age	Car value	Published rides	Posts per month	Reputation	Km
No strike	0.1483	0.7258	28.5598	21.5242	37.5912	6.1926	31.2066	1.6434	4.6398	432.0588
Strike	0.1469	0.7291	28.4609	22.3024	38.0868	6.1491	31.4207	1.6502	4.6397	426.7940

Note: Means of selected variables

variable Table2.15.

Table 2.15: Difference in differences estimation sold seats as a dependent variable

	Dependent variable:		
	sold seats		
	(1)	(2)	(3)
treated	-0.042^{***} (0.013)	-0.024 (0.015)	-0.022 (0.015)
after	-0.154 (0.135)	-0.151 (0.148)	-0.163 (0.148)
did	0.062^{***} (0.023)	0.050^{*} (0.026)	0.050^{*} (0.026)
minority	-0.012^{***} (0.003)	-0.006^{*} (0.004)	-0.003 (0.004)
Driver characteristics			x
Listing characteristics		x	x
Route effects	x	x	x
Time effects	x	x	x
Observations	300,636	243,407	243,407
R^2	0.032	0.033	0.035
Note:			*p<0.1; **p<0.05; ***p<0.01

Alternative definition of treated: minority drivers with less than three reviews driving on the day of strike (table 2.16). We see a higher significance on the treated status.

Proofs

Proof of Proposition 5.

Proof. To find the optimal effort schedule, we first note that efforts exerted at period t does not influence profits in period t, nor previous periods. It does affect future profits, starting from $t+1$. Profit-maximizing drivers will chose efforts at time t such that the marginal cost of efforts equates marginal profits,

Table 2.16: Difference in differences estimation with revenue as dependent variable

	Dependent variable: revenue		
	(1)	(2)	(3)
treated	−0.765 (0.539)	−0.831 (0.726)	−0.925 (0.725)
after	−3.594 (3.671)	−3.805 (4.020)	−3.998 (4.018)
did	0.952* (0.577)	1.304** (0.620)	1.340** (0.620)
minority	−0.556*** (0.085)	−0.382*** (0.097)	−0.316*** (0.097)
Driver characteristics			x
Listing characteristics		x	x
Route effects	x	x	x
Time effects	x	x	x
Observations	297,006	240,473	240,473
R^2	0.040	0.042	0.043

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

Note:Treated: minority drivers with less than 3 reviews

and relation (2.8) simplifies to:

$$g'(a_{it}^*) = \sum_{k=t+1}^{\infty} \beta^{k-t} \mathbf{E}\left[\frac{\partial \pi_{ik}}{\partial a_{it}}\right]$$

First, we calculate the derivative of per-period profits (equation 2.9) at $k > t$ with respect to effort at t:

$$\begin{aligned}\frac{\partial \pi_{ik}}{\partial a_{it}} =& M_k \cdot \left(\left(\frac{\partial S_{ik}}{\partial w_{ik}}\frac{\partial w_{ik}}{\partial a_{it}} + \frac{\partial S_{ik}}{\partial p_{ik}}\frac{\partial p_{ik}}{\partial w_{ik}}\frac{\partial w_{ik}}{\partial a_{i,t}}\right)(p_{ik}^* - c_i) + \frac{\partial p_{ik}}{\partial w_{ik}}\frac{\partial w_{ik}}{\partial a_{i,t}} S_{it}\right) \\ =& M_k \frac{\partial S_{ik}}{\partial w_{ik}}\frac{\partial w_{ik}}{\partial a_{it}}(p_{ik}^* - c_i) \end{aligned} \tag{2.20}$$

, where the second equality stems from driver's price optimization. From the expression of the market share in equation (2.10) we derive the elasticity of the demand with respect to w_{ik}:

$$\frac{\partial s_{ik}}{\partial w_{ik}} = \alpha s_{ik}(1 - s_{ik}) \tag{2.21}$$

Observe that the marginal effect of effort on perceived output for $k > t$ simplifies to:

$$\frac{\partial w_{ik}}{\partial a_{it}} = \frac{\partial}{\partial a_t} \left\{ \frac{h_1 m_1}{h_{mk}} + \frac{h_\epsilon}{h_{mk}} \sum_{s=1}^{k-1} \left(m_1 + a_s - \mathbf{E}\left[a_{is}^*(y^{s-1}) \right] + \mathbf{E}\left[a_k^*(y^{k-1}) \right] \right) \right\} = \frac{h_\epsilon}{h_{mk}} \tag{2.22}$$

where $h_{mk} = h_m + (k-1)h_\epsilon$. Inserting results (2.21) an (2.22) into (2.20), we obtain:

$$\frac{\partial \mathbb{E}(\pi_{i,k})}{\partial a_{i,t}} = \mathbb{E}[M_k]\alpha \frac{h_\epsilon}{h_{mk}}(p_{ik}^* - c_i) \tag{2.23}$$

Assuming that drivers set Bertrand prices:

$$p_{i,k}^* = c_i + \frac{s_{i,k}}{\frac{\partial s_{i,k}}{\partial p_{i,k}}} \tag{2.24}$$

, and noting that the elasticity of demand with respect to price is given by:

$$\frac{\partial s_{ik}}{\partial p_{ik}} = -\gamma s_{ik}(1 - s_{ik})$$

,we obtain the expression for the optimal level of effort with discrete choice demand and Bertrand pricing:

$$\frac{\partial \mathbb{E}(\pi_{i,k})}{\partial a_{i,t}} = \frac{h_\epsilon}{h_{mk}} \frac{\alpha}{\gamma} \mathbb{E}[M_k s_{ik}]. \tag{2.25}$$

All terms of this expression are bounded. Further, we observe that $\sum_{s=0}^{n} \beta^s$ is a converging sequence when $|\beta| < 1$. Hence, it is a classical result that $\sum_{s=t}^{+\infty} \beta^s$ converges to 0 as t goes to infinity. From this, we derive the first part of the proposition.

We finally observe that, when $\forall k \in \mathbb{N}, \mathbf{E}\left[M_{ik} s_{ik}\right] = Q_i$:

$$g'(a_{it+1}^*) - g'(a_{it}^*) = Q_i \frac{h_\epsilon \alpha}{\gamma} \sum_{k=1}^{+\infty} \beta^k \left(\frac{1}{h_{t+k+1}} - \frac{1}{h_{t+k}} \right) < 0 \tag{2.26}$$

h_t being an increasing sequence this expression is negative, which completes the proof. $\square$

Proof of corollary 1:

Proof. The proof follows directly from observing that the expected market share increases as beliefs about the type are revised upwards. The initial belief about the quality μ_m is lower than the true mean $\hat{\mu}_m$, the posterior at time t is given by:

$$\forall a^{t-1}\ \mathbb{E}\left[w_t(\mu_m, \hat{\mu}_m, a^{t-1})\right] = \frac{h_m \mu_m}{h_{m,t}} + \frac{h_\epsilon}{h_{m,t}} \sum_{s=1}^{t-1} \beta^{t-s}(\hat{\mu}_m + a_s - \mathbb{E}\left[a_s^*(w^{s-1})\right] + \mathbb{E}\left[a_t^*(w^{s-1})\right] > \tag{2.27}$$

$$\frac{h_m \mu_m}{h_{m,t}} + \frac{h_\epsilon}{h_{m,t}} \sum_{s=1}^{t-1} \beta^{t-s}(\mu_m + a_s - \mathbb{E}\left[a_s^*(w^{s-1})\right] + \mathbb{E}\left[a_t^*(w^{s-1})\right] = \mathbb{E}\left[w_t(\mu_m, \mu_m, a^{t-1})\right]$$

,which implies that at any $t > 0$

$$\mathbb{E}\left[s_t(\mu_m, \hat{\mu}_m)\right] > \mathbb{E}\left[s_t(\mu_m, \mu_m)\right]$$

Higher market shares imply higher marginal return from effort:

$$\sum_{k=t+1}^{\infty} \beta^{k-t} \frac{h_\epsilon}{h_{mk}} \frac{\alpha}{\gamma} \mathbb{E}[M_k s_{ik}(\mu_m, \hat{\mu}_m)] > \sum_{k=t+1}^{\infty} \beta^{k-t} \frac{h_\epsilon}{h_{mk}} \frac{\alpha}{\gamma} \mathbb{E}[M_k s_{ik}(\mu_m, \mu_m)]$$

Optimal level of efforts equates the marignal cost of providing effort with returns from it; higher returns imply higher optimal level of effort. $\square$

Grades do not depend on prices

We investigate whether grades depend on the prices. We regress price, reputation, plus a full set of other controls on grades obtained. We find that in the OLS estimation there is a positive impact of prices on grades. However, after instrumenting the prices with cost shocks and controlling for driver-specific unobservable effect, we find that the effect is statistically insignificant.

Table 2.17: Impact of prices on grades

	Dependent variable:	
	grade	
	OLS	*panel IV*
	(1)	(3)
price	0.003* (0.002)	−0.016 (0.067)
reputation	0.655*** (0.021)	0.483*** (0.076)
Observations	10,828	1,072
Driver FE		x
Driver characteristics	x	x
Time effects	x	x
Route effects	x	x
Listing effects	x	x
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: regression (1) OLS pooling estimatiot. Regression (2) within driver variation in prices. Prices instrumented with cost shocks: time and space variation in prices and highway tolls

Random coefficients demand estimation

We assumed that the utility of passengers is fully captured by drivers' observed characteristics and a random component. We can thus form *driver categories* that are demand relevant and can be useful for our inquiry: we divide drivers into categories based on the number of reviews: 0, 1-2, 3-4,5-9, and more than 10, together with a minority status (so a category is, for example, zero reviews and not a minority). We aggregate market shares into these categories: thus assuming that passengers are indifferent between any driver in a category. We use these categories as product IDs in a classical BLP setting; this approach has a valuable feature of mitigating the problem of zero market shares. However, we still have some markets where not all product categories are present. We also introduce a random component on price coefficient. Thus, our demand specification takes the following form:

$$Q_{i,t} = M_t \int \frac{\exp(\alpha \bar{w}_{i,t} + \xi_i + \gamma_{j,t} \bar{p_{i,t}} + \phi_t)}{1 + \sum_k \exp(\alpha \bar{w}_{k,t} + \xi_k + \gamma \bar{p_{k,t}} + \phi_t)} dH(\gamma_{j,t})$$

,where $\bar{w}_{i,t}$ is the average price within a category of drivers, ξ_i is a driver category dummy, and H is the join distribution of passenger heterogeneity in $\gamma_{j,t}$.

To address the standard problem of the endogeneity of price, we employ two instrumental strategies. First, we use cost-shifters: over time, the price of gas changes, and we can observe the average price at gas stations in any given city on any given day.[40] These prices change over time (because of oil price fluctuations) and location (e.g., due to varying intensities of competition between filling stations). Additionally, the level of highway tolls varies across routes. Second, we observe the characteristics of all drivers available in a given market: we derive measures of isolation in characteristics spaces.

There are many small markets in our dataset; we have more than 64000 markets, with sometimes fewer than five drivers per market. Therefore, we often observe zero market shares. As noted by Gandhi et al. (2013), a typical "fix" in such a case is to add a small ϵ to all market shares or drop observations with zero market share, which effectively lumps them with the outside option. Unfortunately, both methods lead to biased estimates. In the baseline model, we add ϵ to the market shares of all categories. Furthermore, in some categories are missing from some markets, which can be correlated with a trip fixed effects: for example more minority drivers on specific routes.

We use Python implementation by Conlon and Gortmaker (2019). Figure 2.24 shows estimated elasticities with respect to price and quality measure.

Figures 2.24c and 2.24d use Reynaert and Verboven (2014) to reweigh instruments. Introducing random coefficient on price does not have a big impact on the magnitude of price elasticity. However, we see that elasticity of demand with respect to price is significantly reduced following the optimal instruments procedure. We conclude that elasticity of price is much higher than that of the quality and that the baseline (standard logit) estimates give a reasonable approximation of the more complex model.

[40] www.prix-carburants.gouv.fr

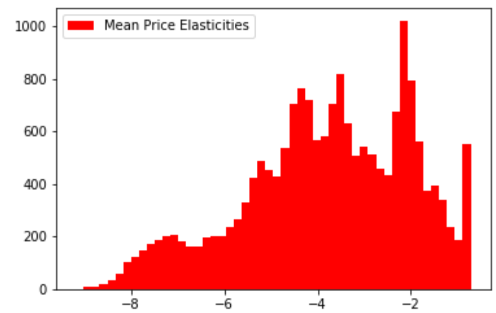

(a) Price elasticity

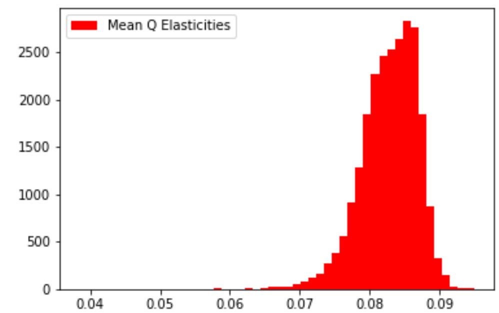

(b) Quality elasticity

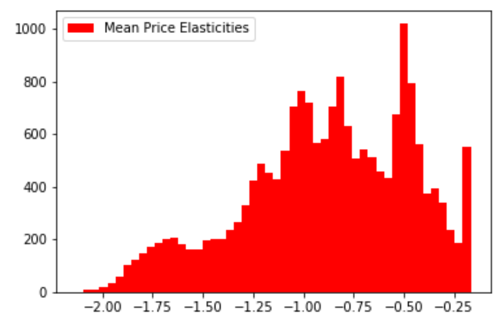

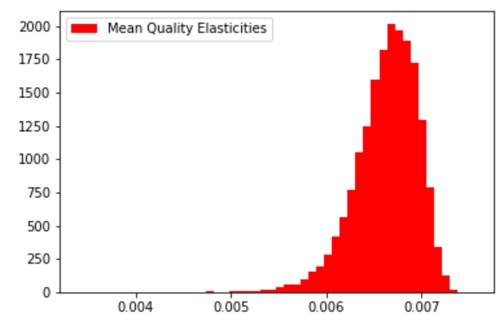

(d) Quality elasticity with optimal instru-

(c) Price elasticity with optimal instrumentsments

Figure 2.24: Random coefficients logit demand.

Demand estimation results all variables

Belief updating with discrete reports

The market forms a prior based on driver's characteristics which are observed on her profile, later on as market receives signals about the performance of the driver, beliefs are updated. Holmström (1999) assumes the prior to be normally distributed with mean and variance: $\eta \sim N(m_1, h_1)$; also, he assumes that signals are distributed normally and continously. This leads to a

	Model 1	Model 2	Model 3
Ride price	$-0.00\ (0.00)^{***}$	$-0.00\ (0.00)^{***}$	$-0.00\ (0.00)^{***}$
Type	$0.12\ (0.06)^{*}$	$0.13\ (0.06)^{*}$	
Log(number reviews)	$0.15\ (0.02)^{***}$	$0.14\ (0.02)^{***}$	$0.15\ (0.02)^{***}$
Automatic acceptance	$0.38\ (0.04)^{***}$	$0.38\ (0.04)^{***}$	$0.39\ (0.04)^{***}$
Picture	$0.56\ (0.22)^{*}$	$0.63\ (0.23)^{**}$	$0.63\ (0.23)^{**}$
Max 2 passengers	$-0.19\ (0.04)^{***}$	$-0.20\ (0.04)^{***}$	$-0.21\ (0.04)^{***}$
Rush time	$0.21\ (0.06)^{***}$	$0.24\ (0.06)^{***}$	$0.24\ (0.06)^{***}$
Day (no rush)	0.08 (0.07)	0.13 (0.07)	0.12 (0.07)
Posted since	$0.06\ (0.00)^{***}$	$0.06\ (0.00)^{***}$	$0.06\ (0.00)^{***}$
Notice	$-0.05\ (0.00)^{***}$	$-0.05\ (0.00)^{***}$	$-0.05\ (0.00)^{***}$
Seniority months	−0.00 (0.00)	−0.00 (0.00)	−0.00 (0.00)
Weekend	$-0.13\ (0.04)^{**}$	$-0.12\ (0.04)^{**}$	$-0.12\ (0.04)^{**}$
Car price		0.00 (0.00)	0.00 (0.00)
Minority		0.06 (0.05)	
Driver Age		−0.00 (0.00)	
Reputation			$0.24\ (0.10)^{*}$
Time effects	x	x	x
Route effects	x	x	x
AIC	31929.66	30150.03	30145.51
R^2	0.45	0.45	0.45
Max. R^2	0.49	0.49	0.49
Num. events	154259	147905	147905
Num. obs.	470165	442839	442839
Missings	0	0	0

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 2.18: Demand estimates: subset of markets.

formation of posterior beliefs:

$$\mathbf{E}\left[\eta|z_s\right] = \frac{h_1 m_1 + h_\epsilon \sum_{s_1}^{t} z_s}{h_1 + t h_\epsilon} \tag{2.28}$$

However, we cannot apply this formula directly because the evaluations are not continuous. Suppose that realizations of output are continuous, but the signals received by the market are discrete. However, there is an objective rule, such that if a realizations falls within a given interval there is always the same grade given: for example, a grade 3 is given when the observed realized output falls within the interval 2.5-3.5, a grade of 5 is given when the observed output is above 4.5. This allows us to calculate marginal probabilities, and characterize the posterior belief, so:

$$\pi(\theta|y) = \frac{f_{y|\theta}(y|\theta)\pi(\theta)}{\int_\Theta f_{y|\theta}(y|\theta)\pi(\theta)d(\theta)} \equiv f_{y|\theta}(y|\theta)\pi(\theta) \tag{2.29}$$

,where $\pi(\theta|y)$ denotes a probability of being of type θ while getting a grade y and $f_{y|\theta}(y|\theta)$ is a conditional probability of a conditional distribution, the empirical counteraprt of equation (4) is

$$\mathbf{E}\left[\theta|Y = y\right] = \frac{P(Y = y \text{ and } \eta = \theta)}{P(Y = y)} * m_i$$

We are currently improving our estimates to account for this.

Estimation of the cost of effort function

We are interested in estimating function $g(a_{i,t})$ that measures the cost of exerting effort. The optimal levels of effort, in our model, are determined by a following relation:

$$a_{imt} = \gamma \left(\sum_{s=t}^{n} \beta^{s-t} \frac{h_\epsilon}{h_{mk}} \frac{\alpha}{\gamma} \mathbb{E}[M_k s_{ik}] \right) + \varepsilon_{ijt} \tag{2.30}$$

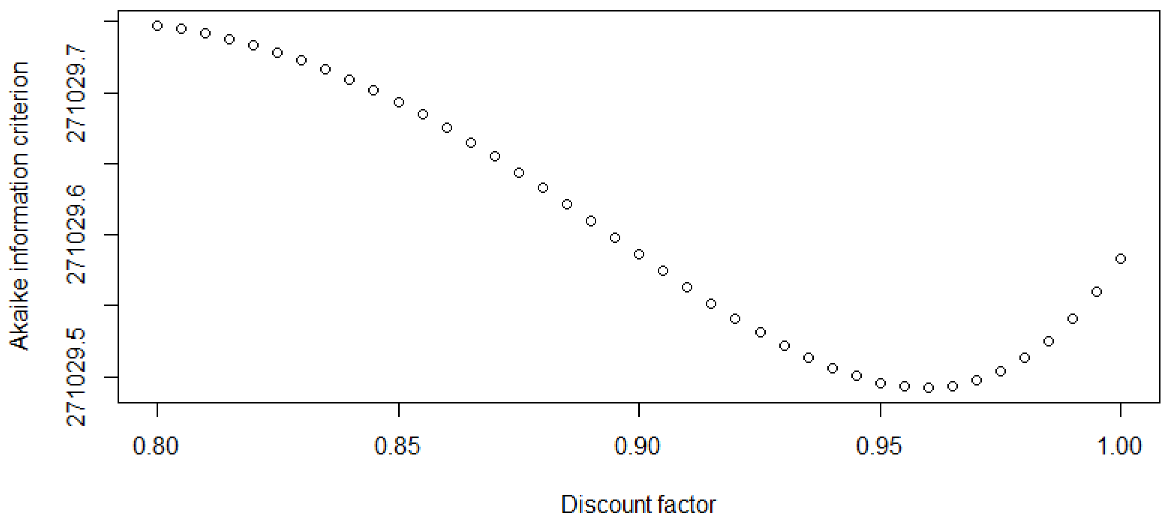

Figure 2.25: Comparison of the quadratic regression of the cost of effort using different discount factors

,where $\gamma(\cdot) = g^{-1,'}(\cdot)$ in the baseline case cost of effort follows a quadratic function. The discounted sum of profits depends on the discount factor. We estimate the quadratic model for different levels of the discount factor and compare the fit using AIC. Figure 2.25 shows the AIC on different levels of the discount factor.

The lowest AIC is achieved for the discount factor of 0.96, and the rest of the models are estimated using this discount factor. In the next steps, we fit higher order Chebyshev polynomials on the discounted sum of profits to compare the fit with the quadratic function. Table 2.19 presents estimates of coefficients of these polynomials

In the next step we present ANOVA results in Table 2.20 From the ANOVA analysis we conclude that inclusion of 2nd and 3rd degree terms improves fit of the model. Finally we present predictions with confidence intervals for linear, quadratic and 3rd model, see Figure 2.26. The difference between the quadratic model the 3rd degree polynomial is clear at the high levels of the

	Dependent variable:				
	effort				
	(1)	(2)	(3)	(4)	(5)
linear	0.97***	−0.62	−12.30***	−14.45	156.12
	(0.05)	(0.47)	(3.17)	(18.97)	(125.35)
2nd degree		4.36***	−12.30***	68.38***	−1,762.95
		(1.29)	(17.23)	(152.66)	(1,351.61)
3rd degree			−113.68***	−174.95	9,656.81
			(30.51)	(535.11)	(18,653.10)
4th degree				79.15	−25,581.11
				(690.23)	(18,653.10)
5th degree					26,304.71
					(19,108.47)
Constant	−0.16***	−0.02	0.67***	0.77	−5.41
	(0.01)	(0.04)	(0.19)	(0.87)	(4.57)
Observations	138,390	138,390	138,390	138,390	138,390
R^2	0.003	0.003	0.003	0.003	0.003
Adjusted R^2	0.003	0.003	0.003	0.003	0.003
Residual Std. Error	0.64 (df = 138388)	0.64 (df = 138387)	0.64 (df = 138386)	0.64 (df = 138385)	0.64 (df = 138384)
F Statistic	426.52*** (df = 1; 138388)	218.98*** (df = 2; 138387)	150.63*** (df = 3; 138386)	112.98*** (df = 4; 138385)	90.76*** (df = 5; 138384)
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 2.19: Estimates of $g(\cdot)$: Column 1 linear model, columns 2-5 polynomials of increasing degrees

Degree	Res.Df	Sum of Sq	F	Pr(>F)	
1	138388				
2	138387	4.7344	11.4098	0.0007308	***
3	138386	5.7617	13.8858	0.0001943	***
4	138385	0.0055	0.0132	0.9086991	
5	138384	0.7863	1.8950	0.1686383	

Table 2.20: ANOVA analysis of models from linear to 5th degree polynomial

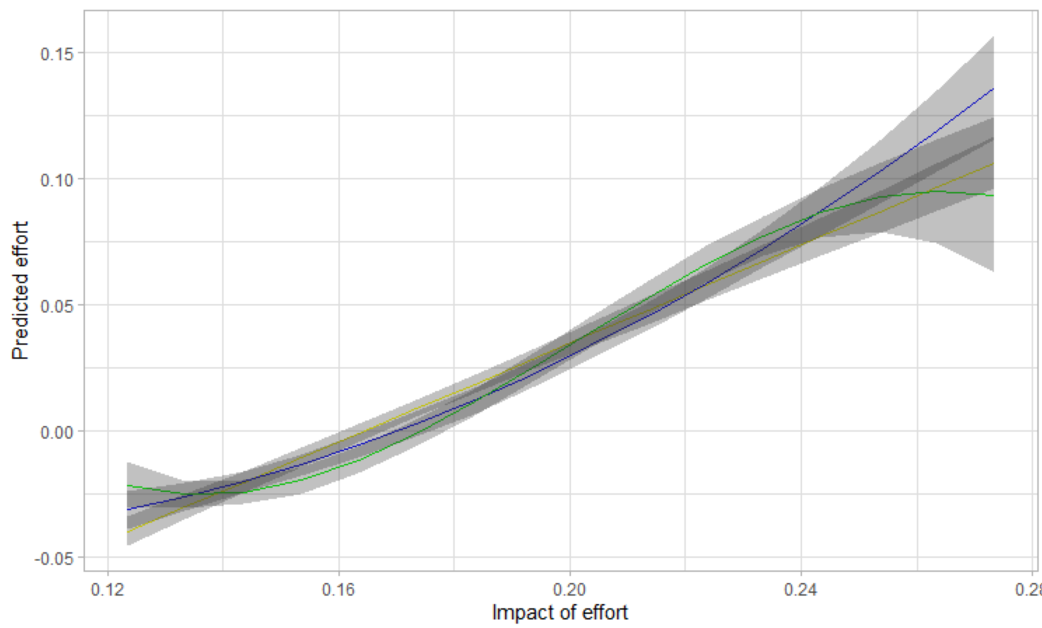

Figure 2.26: Comparisons of predicted effort: Yellow linear model, Blue- quadratic model, Green- 3rd degree polynomial

horizontal axis.

Chapter 3

The price is right! Information and dynamics in online marketplaces

This chapter is based on a research project with Rossi Abi Rafeh.

3.1 Introduction

The internet transformed many industries by enabling the entry of small often part-time sellers and allowing them to compete against traditional players: Airbnb facilitated short-term house rentals creating a market in which home-owners compete with hotels; Uber matches passengers to drivers for short-distance trips, competing with taxis; Blablacar connects drivers to passengers on city-to-city trips, competing with airlines, trains, and buses; Upwork connects companies to freelancers, competing with interim agencies.

These platforms created marketplaces for assets and capital that were previously under-utilized: before Blablacar, drivers from Toulouse to Paris had to either pick up hitchhikers from the highway toll zone or travel with empty seats (or take a train themselves). As these sellers are often part-time workers, they have to decide how often to enter the market. These entry decisions determine the depth of the seller's side of the market and the attractiveness of the platform for the buyers.

Despite profound differences in platform-deisgn[1] the high turnover rate of sellers is a common feature. According to a study by JP Morgan Chase (JP Morgan Chase & Co Institute (2018)), 56% of people offering their capital to rent on sharing economy platforms (Airbnb, etc...) quit within their first year, 52% of those offering their labor services on a sharing platform do as well. Understanding the pricing and entry/exit decisions of sellers is key to predict which design decisions are efficient, or optimal to the platform.

In this paper, we study the determinants of sellers' entry and pricing behavior in an online marketplace. We emphasize the importance of unobserved (by econometricians) sellers' heterogeneity in marginal and opportunity costs. The marginal cost is the disutility a seller has in providing the service to an additional buyer: on Airbnb, this is the effort to leave one's apartment to strangers and preparing it for them; for ride-sharing, this is the disutility of going to a pick-up location or having a stranger in the car during the trip.

[1] e.g. Uber sets the price of the ride, and automatically matches passengers to drivers; Blablacar, on the other hand, leaves the pricing decision to the driver, and the choice of the match to the passenger.

While the opportunity cost is associated with foregoing some other activity; on Blablacar this could be refraining from traveling by train or even the cost of committing to time and place of departure. These costs are unrelated to the number of buyers served, but necessary to provide the service. We argue that accounting for this cost structure is important to give predictions in the equilibrium on the platform: entry, exit as well as pricing equilibrium of sellers.

The presence of reputation systems on online marketplaces introduces a static-dynamic trade-off in sellers behavior: a seller would enter the market if she can set a high enough price; however, lowering prices in early listings can increase reputation and boost future profits. To account for this mechanism we construct a dynamic model of oligopolistic competition à la Ericson and Pakes (1995). Sellers compete in prices and make decisions to enter the market in each period. Their reputation levels develop following a probabilistic process and act as state variables.

To validate the predictions of our model and to highlight the importance of the unobserved costs we use data on a large ride-sharing platform from Lambin & Palikot (2019). We start by showcasing an interesting empirical puzzle: in the cross-section, entrant sellers set higher prices than incumbents, despite the accumulated reputation. After controlling for a rich set of driver and listing-specific controls we find that the average effect of an extra review on the price is negative and significant. Second, we match listings posted by the same seller and study within driver-variation, we recover the expected price dynamic: sellers start at a low price in their early listings on the platform and increase it gradually. Both the first-difference (FD) and within (FE) estimators show a positive and significant impact of reputation on prices: going from 0 to 10 reviews now *increases* the average price by 70 cents. The Hausman test rejects the uncorrelatedness of the driver-specific unobservable effects with the number of reviews and the other observables.

Second, we estimate a structural demand model and find that consumers are price-sensitive, consistent with findings from the literature on online marketplaces. Passengers are more responsive to the depth of a driver's reputation (number of reviews) than the average quality of reviews (number of stars).

Assuming Bertrand-Nash in pricing (for experienced drivers), we derive the marginal costs and observe a substantial heterogeneity: some drivers derive positive utility of sharing a ride with a stranger while most do not and price to be compensated accordingly. Using the estimated costs, we show that the likelihood of exit from the platform, defined as not having posted any new listing over six months is highly correlated with the estimated marginal costs giving credence to a mechanism of selection on marginal costs. This effect holds for controlling for the quality and quantity of seats sold.

Finally, we simulate the supply-side model using the demand-side parameters. We show that sellers with a high marginal cost set higher prices; however, they do not enter the platform when they receive a high draw of the opportunity cost. Sellers with low marginal costs remain on the platform and start building a reputation, which allows them to start gradually increasing the prices. Relatively high heterogeneity in marginal costs is key to generate such a price dispersion.

The paper proceeds as follows. In section 2 we discuss related literature, section 3 introduces empirical our model. Section 4 discusses the empirical test of our predictions. Section 5 shows the results from a structural model and section 6 concludes.

3.2 Literature review

This paper tackles the problem of optimal dynamic pricing by firms competing on a platform, in the presence of incomplete information about their quality. Hence, it falls into several strands of economic literature. First, pricing by long-lived economic agents has been studied by both theoretical and empirical literature. Maskin and Tirole (1988a,b) have introduced the equilibrium concept on which most of the literature on dynamic oligopoly is based. These papers describe features of a Markov Perfect Equilibrium (MPE) and derive dynamic programming equations that provide a solution for an MPE. Jovanovic (1982) studies the evolution of a market structure focusing on the problem of selection. In this work, firms learn about their efficiency as they operate, and exit when they realize that their marginal costs are too high. This dynamic is also exhibited in our results.

Numerical solutions to an MPE in which firms invest in R&D are introduced in Pakes and McGuire (1992) and Ericson and Pakes (1995). Our simulations draw heavily from their algorithm. Their framework has been extended by many authors in the last 25 years, Doraszelski and Pakes (2007) provide an overview of the main contributions. Recently several papers have advanced the idea of prices as investments (see for example Besanko et al. (2014, 2010)). They develop a model with a learning-by-doing feature and showcase firms' incentives to strategically decrease their prices to advance their technology, but also to deny technological improvements to competitors. These papers focus on a limited number of competitors (typically two or three) that compete in each turn; such a restriction is due to the computational complexity of the MPE. However, a handful of firms competing against each other repeatedly does not reflect markets that we want to study. Often sellers compete with a potentially large and uncertain number of players, whose identity changes from period to period. Weintraub et al. (2008, 2010) introduce a concept of an Oblivious Equilibrium which is shown to be a good approximation of an MPE, despite being computationally much less demanding.

Second, previously mentioned papers have explored R&D investments and a learning-by-doing mechanism to model industry evolution, our work focuses on the reputation of sellers as a state variable. The literature on the value and the impact of online reputation has emerged at the beginning of the 2000s, and many early papers were focused on eBay's reputation system. Livingston (2005) shows that sellers on eBay are strongly rewarded with the first few positive reviews; Cabral and Hortacsu (2010) also confirm these results, they also show that early negative reviews are of particular importance. Jolivet et al. (2016) study the impact of seller's reputation on equilibrium prices; notably, by close attention to unobserved sellers heterogeneity, their approach is related to our reduced form results. Reputation has also been shown to matter in an offline set-up by for example Spagnolo (2012). This paper is particularly important for us because it models entry to the market with a reputation system.

Finally, we assume that selles have a quality component that is initially unknown both to them and to the market. Over time, information about the sellers is revealed through reviews left by the passengers. Thus, we model sellers as experience goods (Nelson, 1970). Therefore, our paper is related to the literature on the dynamic pricing of experience goods. Bergemann and Välimäki (2006) analyze the problem of a monopolist selling a new experience good without any initial private information about its quality. They show that depending on the discount factor two price profiles can emerge: in "niche markets", the monopolist starts with low introductory prices and increases over time; in "mass markets" the opposite price profile arises. Bergemann and Välimäki (1996, 2000) analyze the interplay between learning and pricing in markest with both entrant and incumbent sellers. The entrant has no private information about its quality. Bergemann and Välimäki (1996) takes the case of a single long-lived buyer. If the entrant makes a sale in the first period, the revelation of information will induce higher prices, in expectation, in later periods. Thus, in all the "cautious" MPEs of the game, the entrant sets introductory prices that are lower than the marginal cost. The interplay between the increased competition for future market shares and the decreased price elasticity of demand of long-lived buyers is revisited in Miguel Villas-Boas (2006). In Bergemann and Välimäki (2000), buyers are short-lived thus very elastic to prices, the entrant makes even deeper discounts in the early periods in order to build a reputation. These papers focus on the level of experimentation that arises in the equilibrium: dynamic pricing leads to an efficient level of experimentation when the buyers are long-lived, but to too much experimentation when buyers are myopic. Hagiu and Wright (2018) extend these results to the context of a platform. Overall, when sellers with no private information compete, the theory predicts, that entrants post low prices and increasing them in subsequent periods.

3.3 Empirical model

We propose a model of a dynamic oligopoly, where the sellers are long-lived, but they have to decide whether to trade or not in each period. If they decide to trade they have to entail an opportunity cost, but they benefit from per-period profit and improve their continuation value by building reputation. Formally, the model is similar to Ericson and Pakes (1995), where all incumbents exit every period, and entry decisions are effective in the same period as when they are taken. Each period t we split into entry decisions, profits, and reputation updating stages. Our empirical strategy follows closely the literature on Experienced Based Equilibria (Fershtman and Pakes, 2012; Pakes et al., 2007).

3.3.1 Demand for ride-sharing services

We assume that passengers looking for a ride from city A to city B compare the drivers available on a given day, and choose the one that maximizes their utility. Demand is then a standard discrete choice setup: passenger i chooses the ride j among all the listings available on route r at day t. The utility of passenger i choosing the seller j writes

$$u_{i,j,r,t} = \alpha_1 p_{j,r,t} + \alpha_2 \mathbb{E}[q_{j,r,t}|I_{j,t}] + \alpha_3 n_{j,r,t} + \alpha_4 nl_{j,r,t} + \beta X_{j,r,t} + \epsilon_{i,j,r,t}, \quad (3.1)$$

where $p_{j,r,t}$ is the price set by seller j when listing on route r and day t, $\mathbb{E}[q_{j,r,t}|I_{j,t}]$ is the perceived quality of seller j given the signal I_{jt} they have received up to time t, $n_{j,r,t}$ is a concave transformation of the number of reviews, and $nl_{j,r,t}$ is a concave transformation of the number of trips the driver has already made on the platform. $X_{j,r,t}$ is a collection of listing-specific variables, observable to both the passenger choosing a listing, and the econometrician. $\epsilon_{i,j,r,t}$ is i.i.d across passengers, sellers, routes and days. Passengers can choose an outside option and get the utility $u_{0,r,t} = \epsilon_{i,0,r,t}$. We allow the outside option to vary by route and day. Conceptually, the outside option includes alternative modes of transport like train or a choice of travel by ride-sharing on a different day.

One might be concerned with the endogeneity of price. In our application, we have individual-level data, although per se this does not alleviate concerns of price endogeneity; however, combined with the availability of all ride-specific characteristics that are observed by the passenger at the time of the choice it reduces the concerns for an omitted variable bias. Nevertheless, the textual description of the listing and content of the photo (beyond age, gender, or ethnicity), and written comments by past passengers are potential sources of omitted variable bias.

3.3.2 Theoretical model of market entry

A large number M of long-lived potential sellers decides every period whether to trade or not. Each seller j is characterized by $x_{j,t} = (n_g, n_b, c)$ at time t. These characteristics include the reputation of the seller (number of good and bad signals they received up to the current period) and their marginal cost of providing the service c_j, all of them being non-negative, discrete numbers with upper bounds at $\overline{n}_g, \overline{n}_b, \overline{c}_j$; the state space Ω takes all potential combinations of $(n_{g,j,t}, n_{b,j,t}, c_j)$. Let Ξ_t denote the number of potential sellers, who decide to trade in the period t; X_t is a vector that summarizes the number of sellers, who trade in period t, and their characteristics; so, for example, the first entry is the number of sellers with no reviews (neither good or bad) and zero marginal cost. Let z_t be a vector of exogenous profit shifters, which evolves following a finite-state Markov process. We assume that per period profits of seller j are fully characterized by these variables: $\pi_{j,t}(x_{j,t}, X_{j,t}, z_t; \theta)$, where θ is a parameter vector, the true value of which is θ_0. The reputation of sellers is updated after profits are realized and it follows a probabilistic process; per period profits are increasing in the number of good reviews and decreasing in the number of bad reviews. Furthermore, we assume no strategic behavior by sellers after the entry decisions are realized. Finally, the market structure X_t is determined by the profit function $\pi(\cdot)$ and the profit shifter z_t; thus, for each value of z_t a potential seller can form an expectation of the number and characteristics of potential competitors.

Each of the potential sellers observes z_t and decides to be active if the expected profits in period t plus change in the value of continuation is higher then the opportunity cost. The opportunity cost is $\phi_{j,t}$ and δ is the discount rate. Sellers problem is described by a Bellman equation,

$$V(x, X, z\phi; \theta) = \max\{\phi + \delta VC(x, z; \theta), \pi(x, X, z; \theta) + \delta VC(x', z; \theta)\}$$

,where x' reflects sellers' characteristics after the expected update to the reputation. Sellers choose to be active when the latter term in the maximization problem is higher. $VC(\cdot)$ is the continutation value, which is the expectation of the next period realization of the value function,

$$VC(x, X, z, \phi; \theta) = \sum_{x',z',X} \int_{\phi'} V(x', X', z', \phi; \theta) p(d\phi'|\theta) p(X'|z, \chi = 1) p(z'|z) p(x'|x)$$

,$p(\cdot)$ denotes the probability distribution: for example $p(x'|x)$ is distribution of sellers state in the next period x' condition on her being in the state x, and $p(X'|z, \chi = 1)$ is the expected market structure in the next period, condition on the market being in state z today and the seller being active next period ($\chi = 1$). Pakes et al. (2007) shows that, in the equilibrium, seller's expectation of the market structure in the future periods has to be consistent with the expectations of others sellers. We use following assumptions to determine the equilibrium behavior.[2]

Assumption 1: The distribution over the opportunity costs $F^{\phi}(\cdot|\theta)$ is non-negative and generates costs i.i.d across time and markets. The distribution is known to all sellers, but the realizations are observed only by a given seller (asymmetric information).

[2]These are very similair to the assumptions in (Fershtman and Pakes, 2012; Pakes et al., 2007).

Assumption 2: Perceptions of the market structure in period $t+1$, $p(X_{t+1}|z_t, \chi = 1)$ is influenced only by the profit shifter z_t. The evolution of the profit shifters follows Markov process, and $\pi(\cdot)$ is bounded.

With these assumptions we characterize an equilibrium, the key observation is that entry decisions, which are simultaneous, are based on the market conditions z that are obseved by everyone and opportunity costs.

Proposition 6. *An equilibrium is a collection of optimal entry decisions of each seller* $\{e^*_{j,t}\}_{t\in 0,\infty}$ *given their opportunity cost and profit shifters* $z_{t\in 0,\infty}$*, and a market structure* $\{X_t\}_{t\in 0,\infty}$ *such that:*

1. *All sellers have consistent perceptions of the likely market structure* $p(X'|z, \chi = 1)$
2. *The entry decision is individually rational for a seller in a state* x_j *with opportunity cost* $\phi_{j,t}$*,* $e^*_{n,t}(\phi, z_t) = \mathbf{1}_{\pi(x_{j,t}, X_t, z_t;\theta)+\delta VC(x',z;\theta)-\phi_{j,t}-\delta VC(x,z;\theta)}$
3. *The market structure* $\{X\}_{t\in 0,\infty}$ *arises with the optimal entry decisions of the sellers.*

Perceptions of the future market structure must reflect that probabilities of entering for sellers of a given state are identical up to their draws on opportunity costs. As long as the distribution of seller types in population M is known sellers can form expectations which will be consistent with equilibrium play. Let the state space Ω be indexed by k, with $K = Card\,[\Omega]$.

$$p(X|z, \chi = 1) = \sum_{k=1}^{K} F^{\phi}\left\{VC(X, z;\theta)|\theta\right\}$$

The crucial observation that will allow us to write down the empirical counterpart of the value functions characterized above is that there is a set of recurrent market structure - profit shifters combinations (X, z), and within each of them the equilibrium play of the sellers gives rise to the same data generating process. If the dataset is large enough, so that all recurrent (X, z)

combinations have been observed, and as a consequence, the realized profits are known as well, we can calculate the continuation values.[3]

3.4 Empirical application

In this section, we apply our model to data from an online marketplace. First, we discuss the dataset that we use. Second, we provide some reduced-form statistics that highlight the importance of unobserved characteristics in understanding market dynamics. Finally, we use the data to estimate a model of demand and calibrate simulations of the supply model.

3.4.1 Data and empirical context

The dataset comes from Lambin and Palikot (2020); we are grateful to the authors for agreeing to share it. Here we briefly present the dataset, please refer to the original paper for the details of the collection, and processing. Tables with summary statistics and the definition of the main variables are in the Appendix 3.7.

Blablacar is a popular French ridesharing platform; it was established in 2006. Today it has around 60 million users worldwide; in France, it serves 1.5 million travelers every month[4]. The platform caters to mostly non-professionals drivers looking to cover the costs of a trip from A to B. Blablacar experienced a period of fast growth in the first ten years of its functioning. However, recent regulatory changes, in particular, liberalization of intercity bus services have increased competition in the intercity transportation market in France, and as a consequence lead to stagnation of Blablacar in France.

There are some critical institutional differences between Blablacar and other ridesharing platforms (e.g., Uber): drivers set their prices, and passengers choose a driver they want to travel with from a list of available alternatives (some drivers reserve a right to reject requests made by passengers).

[3]In an ongoing project with Rossi Abi Rafeh, we are structurally estimating this supply model. In Appendix 3.7 I briefly discuss the estimation strategy.

[4]https://blog.blablacar.fr/about-us/qui-sommes-nous

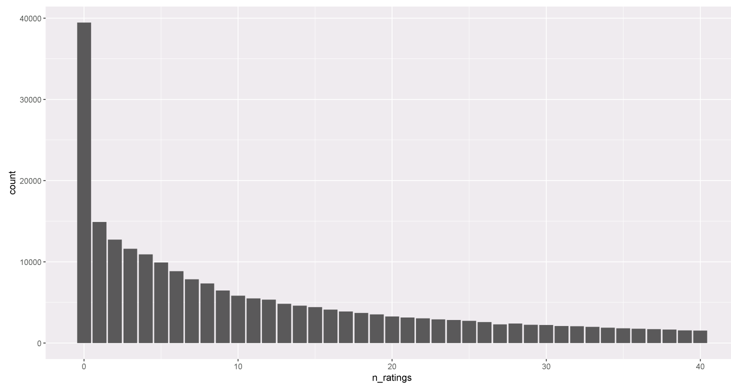

Figure 3.1: Share of drivers on different levels of experience.

The decision to enter the market and propose a ride on a given day reflects the driver's expectation of profits she would make taking into account the opportunity cost of doing that (e.g. committing to a time of the trip or forgoing other means of transportation). These individual decisions give rise to the market structure on a given day. Figure 3.1 shows that typically there are many more entrants to the platform then experienced users; this observation combined with the fact that Blablacar is not a growing marketplace, reveals an interesting observation: most of the entrants leave the platform soon after creating a profile. This observation is consistent with other studies (JP Morgan Chase & Co Institute (2018)).

One might expect that sellers on different levels of experience will price differently. If returns from reputation are sufficiently high prices will increase as drivers gain experience. On the contrary, if there is a significant amount of heterogeneity in costs, the composition of the sample will change as drivers exit: that might lead to more efficient drivers staying, and thus the average price could even decrease. Figure 3.2 plots prices against the number of reviews; this suggests that more experienced drivers set lower prices than entrants.

Figure 3.3 shows the distribution of prices for users with no reputation, some reputation, and experienced ones. We see that prices set by drivers with no reviews are skewed to the left, and this skewness increases as drivers become more experienced. There is a mass of high prices that gradually disappears.

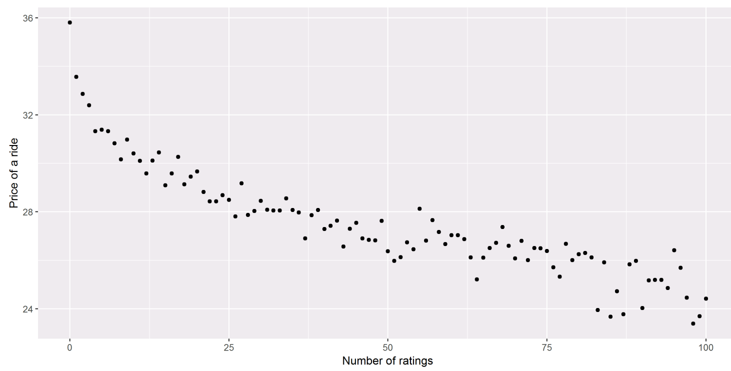

Figure 3.2: Mean prices on different stages of experience.

At first, these patterns are quite surprising, seem to be contradictory with findings of the literature on dynamic pricing of experienced goods (Bergemann and Välimäki, 2000), and with prior empirical work on prices in online markets (Jolivet et al., 2016).

3.4.2 Empirical strategy

In this subsection, we use a reduced-form econometric model to, first, study the effect of early ratings on the pricing of drivers. Second, we relate pricing behavior, and ratings to the probability of exiting the platform.

Let $p_{j,r,t}$ denote the price posted by seller j, on a route r (pair of origin and destination cities) on day t. The econometric specification we use in this section follows the model

$$p_{j,r,t} = g(n_{j,t}) + q_{j,t} + \gamma X_{j,r,t} + \lambda_t + \lambda_r + \mu_j + \epsilon_{j,r,t}, \tag{3.2}$$

where price is a function of $n_{j,t}$ the number of past ratings a driver has accumulated at the time of the ride, $q_{j,t}$ the average of the ratings, $X_{j,r,t}$ time/route/driver characteristics that are observable to the econometrician, and time/route/driver fixed unobservable characteristics; μ_j denotes the driver-specific time-invariant characteristics that are relevant to the price-setting problem of the driver but are unobserved to the econometrician; $\epsilon_{j,r,t}$ are idiosyncratic shocks that affect the pricing of the driver. We assume that

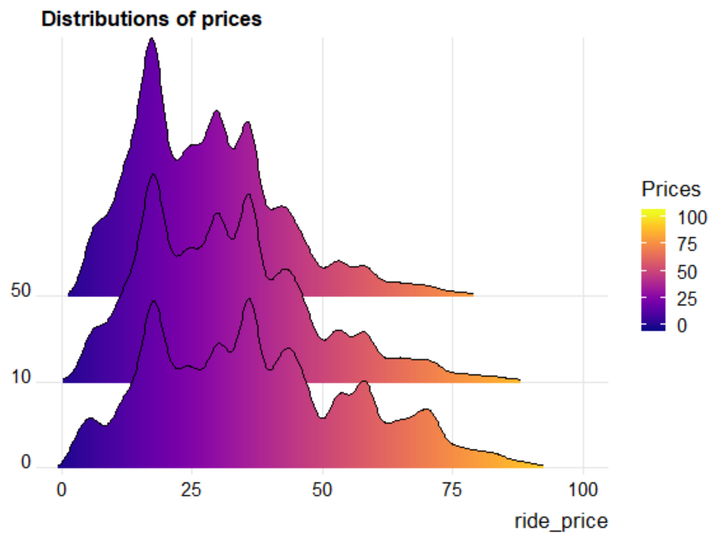

Figure 3.3: Distributions of prices on different stages of experience. Vertical axis number of reviews

the shocks $\epsilon_{j,r,t}$ are independent of the regressors at the time of the ride: $\mathbb{E}(\epsilon_{j,r,t}|q_{j,t}, n_{j,t}, X_{j,r,t}) = 0$

The function g characterizes the dependence of the seller's price on the number of ratings that she has collected before the time t. In practice, we estimate a non-parametric effect of binned ratings. The binning is thinner at early ratings and coarse for more experienced drivers: for example, the first bin consists of listings that have zero ratings, but all ratings above 100 constitute one bin.

Identification Our goal is to identify the effect of early ratings on sellers' prices. Are entrant sellers setting higher prices and decreasing them as they become more experienced, or are they starting with low introductory prices and increasing their prices gradually as they gain more reputation?[5]

Our identification strategy relies on including a rich set of observable char-

[5]There might be other, non-costs, reasons for a seller to set a low introductory price, for example, in chapter two, I show in the same context that sellers from ethnic minorities have an incentive to decrease price upon entering the market in order to counter prejudice that they face.

acteristics and exploiting the panel-like structure of our data. First, we have a measure of the "objective" travel cost of a route for a driver: it is a function of the distance, the fuel consumption of the driver's car, the price of gasoline on the day of the trip in cities of departure and arrival and the highway tolls on the route traveled. The heterogeneity in route-day-specific travel cost, along with the heterogeneity in the driver-specific observable costs account for most of the variation in prices in our results below. We also include driver-specific controls (age, sex, automatic acceptance of the rider, ethnicity, etc.), and route/time-specific controls (city, trip, the day of the week, strike day, etc.) Second, most of the drivers post several trips over their career, a sub-sample of these listings is observed for each driver. Thus, we can use within-driver variation in order to identify the effect of early ratings on the price set.

The primary source of endogeneity that arises in our context is due to unobserved driver-fixed effects that determine both the price that the driver sets and the number of ratings she already has. For instance, drivers with high costs are less likely to post rides, and when they do, they set higher prices than their peers. These unobserved driver fixed effects appear to be the main source of bias in our context.

Another source of endogeneity is the potential correlation between the current number of ratings $n_{j,t}$ and past realizations of the idiosyncratic shock $\epsilon_{j,r,t}$. Although this is not an issue in the model in levels, the first-difference or within estimators which we use in order to exploit the within-driver variation might be biased if $E(\epsilon_{j,r,t}|n_{j,t'}) \neq 0$ for some t, t'. We believe this bias may not be prevalent in our case as the first-difference and within estimates are similar: the bias created by such a correlation would result in a discrepancy between the first-difference and the within estimators. The difference in the estimates results from the fact that the two estimators use distinct methods to eliminate the fixed effect from the regression equation. The FD estimator bias would arise from $E(\epsilon_{j,r,t} - \epsilon_{j,r,t-1}|n_{j,t} - n_{j,t-1}) \neq 0$ whereas the within estimator bias would result from $E(\tilde{\epsilon}_{j,r,t}|\tilde{n}_{j,t}) \neq 0$, where $\tilde{n}_{j,t}$ is the within-driver demeaned number of ratings.

For a large share of drivers, we observe (1) more than two rides; and (2) rides posted at very different points of their careers (early on vs. later on). The effect of a past shock on today's number of ratings is decreasing with time, and one can use long-differences or Arellano-Bond style internal instruments to account for this source of endogeneity. In our analysis, we take the view that ratings are signals of the underlying quality of the driver. The noise in the past signals is independent of the past prices of the transaction, thus of the past idiosyncratic shocks: $\mathbb{E}(\epsilon_{j,r,t}|q_{j,t'}) = 0 \forall t, t'$: If ratings given by a customer depend on prices that they paid, we have a similar problem, which long-differences and internal Arellano-Bond instruments could address.

3.4.3 Results

In this section, we use a subset of rides spanning a year from July 2017 until August 2018, in total there are 302 502 observations. After controlling for the observables, we note that sellers with a higher number of reviews set lower prices than entrant sellers who do not have yet reputation. The average effect is negative and significant: an additional review is associated with a decrease of 3 cents of the euro (1% of the average price observed in the sample). In column (5) of Table 3.1, we report the estimates for the regression using the full sample and binned ratings. The effect is more accentuated at the early ratings: first four ratings reduce prices by 23 cents, moving from 4 to 16 ratings decrease prices by a further 37 cents; finally, ratings from 16 to 40 have a cumulative effect of 17 cents reduction in price. The effect is monotonously decreasing in the number of ratings.

As mentioned before, a unique aspect of the dataset that we use is that we have a useful measure of the objective or physical marginal cost of a ride for the driver. The travel cost is the single most important variable in predicting the price; without any further controls, it explains around 80% of the variation in drivers' prices. From column (3) in Table 3.1, an increase of 1 euro in the travel cost is associated with an increase of 47 cents in the price of a ride on average. Controlling for trip fixed effects takes out the route-specific common cost between drivers, and the much lower estimates in columns (4) and (5) now reflect the effect of between-driver heterogeneity. Drivers with highly efficient cars pass-through only a small fraction of this reduction in cost to their passengers (approx 2 cents for one standard deviation of fuel efficiency).

Many drivers appear in our dataset multiple times at various stages of their careers; hence, we can control for unobserved driver-specific fixed effects. We estimate two transformations of the price regression equation: a first-difference and a within-driver transformation. We include the binned number of ratings as before. In contrast to the OLS results, both estimators show a positive and significant impact of reputation on prices: therefore, once the unobserved driver characteristic is accounted for in the pricing decision, drivers with more ratings set on average higher prices. Collecting the first four ratings now *increases* the average price by 60 cents, going from 4 to 16 by a further 14 cents. The FD and within estimates are similar, and reflect a similar dynamic of pricing by the entrants: entrant sellers with no or few reviews price lower than sellers with an established reputation.

As mentioned before, we expect a non-zero correlation between past shocks to pricing and current number (and quality) of ratings, but this correlation does not give rise to the same bias in the within and FD estimates. We do not formally correct for this bias. We expect doing so would decrease our FD and within estimates of the effect of a rating, but we do not think it will change signs of the coefficients.

Results presented in Table 3.2 indicate that the unobserved heterogeneity plays an important role. We argue that these unobserved time-invariant characteristics are constituted of driver-specific (marginal) costs of providing the ride-sharing service. The drivers with, on average, higher prices will sell fewer seats and in result obtain fewer reviews. Thus, there is a negative relationship between the cost of the driver and the number of reviews.

Furthermore, if the higher marginal cost is correlated with higher opportunity costs, then the unobserved heterogeneity in opportunity costs would also create a negative correlation between the number of ratings and the price, further negatively biasing the OLS estimates.

This is what observe in Table 3.2, the pooling estimator (Column 1) indi-

Table 3.1: OLS on cross-section

	Dependent variable:				
	Price of a ride				
	(1)	(2)	(3)	(4)	(5)
Constant	39.822***	3.979***	1.490***	52.820***	54.200***
	(0.461)	(0.220)	(0.338)	(0.319)	(1.191)
1-2 reviews	−2.453***	−0.165***	0.051	−0.029	−0.121***
	(0.122)	(0.060)	(0.060)	(0.040)	(0.040)
3-4 reviews	−3.734***	−0.479***	−0.232***	−0.154***	−0.231***
	(0.130)	(0.063)	(0.063)	(0.041)	(0.041)
5-8 reviews	−4.779***	−0.709***	−0.485***	−0.309***	−0.378***
	(0.115)	(0.056)	(0.056)	(0.037)	(0.037)
9-12 reviews	−5.451***	−0.806***	−0.555***	−0.404***	−0.469***
	(0.129)	(0.062)	(0.062)	(0.041)	(0.041)
13-16 reviews	−6.029***	−1.084***	−0.812***	−0.577***	−0.607***
	(0.140)	(0.066)	(0.066)	(0.044)	(0.044)
17-20 reviews	−6.116***	−1.078***	−0.785***	−0.546***	−0.569***
	(0.152)	(0.071)	(0.071)	(0.047)	(0.047)
21-40 reviews	−7.599***	−1.580***	−1.262***	−0.807***	−0.773***
	(0.108)	(0.052)	(0.053)	(0.035)	(0.035)
41-100 reviews	−9.405***	−2.265***	−1.897***	−1.113***	−0.961***
	(0.107)	(0.052)	(0.052)	(0.035)	(0.036)
101+ reviews	−12.998***	−2.765***	−2.356***	−1.326***	−1.043***
	(0.117)	(0.057)	(0.057)	(0.038)	(0.039)
Average rating	−0.809***	0.096**	0.060	−0.199***	−0.339***
	(0.099)	(0.046)	(0.046)	(0.030)	(0.030)
Travel cost		0.478***	0.474***	0.022***	0.024***
		(0.0005)	(0.0005)	(0.001)	(0.001)
Travel Cost		X	X	X	X
Time effects			X	X	X
Trip effects				X	X
Driver Controls					X
Observations	298,185	238,509	238,509	238,509	235,082
R^2	0.054	0.832	0.835	0.929	0.931
Adjusted R^2	0.054	0.832	0.835	0.928	0.930

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

Table 3.2: Repeated observation sample

	Dependent variable:		
	Price of a ride		
	Pooling	First-Difference	Within
	(1)	(2)	(3)
Constant	55.994***	−0.016	
	(1.514)	(0.012)	
1-2 reviews	−0.075	0.358***	0.404***
	(0.051)	(0.086)	(0.078)
3-4 reviews	−0.164***	0.531***	0.587***
	(0.052)	(0.090)	(0.079)
5-8 reviews	−0.309***	0.691***	0.738***
	(0.048)	(0.093)	(0.078)
9-12 reviews	−0.375***	0.692***	0.816***
	(0.052)	(0.102)	(0.084)
13-16 reviews	−0.560***	0.710***	0.740***
	(0.055)	(0.111)	(0.090)
17-20 reviews	−0.461***	0.778***	0.763***
	(0.058)	(0.119)	(0.095)
21-40 reviews	−0.721***	0.623***	0.667***
	(0.046)	(0.121)	(0.094)
41-100 reviews	−0.914***	0.338**	0.443***
	(0.045)	(0.136)	(0.107)
101+ reviews	−1.012***	0.133	0.057
	(0.047)	(0.167)	(0.129)
Average rating	−0.399***	−0.165*	−0.143*
	(0.034)	(0.099)	(0.086)
Travel cost	0.025***	0.034***	0.037***
	(0.002)	(0.004)	(0.003)
Travel cost	X	X	X
Time effects	X	X	X
Driver Controls	X	X	X
Trip effects	X	X	X
Observations	165,205	113,131	165,205
R^2	0.935	0.894	0.893
Adjusted R^2	0.935	0.893	0.844
Note:			*p<0.1; **p<0.05; ***p<0.01

cates a negative effect of the number of early ratings on price, but accounting for the unobserved heterogeneity, the effect of early ratings on price becomes positive (Column 2 and 3).

Second, we provide another robustness check by restricting the dataset to drivers that stayed on the platform long enough to receive 40 reviews, and that were observed in the dataset when they had less than five reviews. This can be seen as selecting the subsample of drivers with low marginal costs: in this subsample, the bias of OLS disappears, which is consistent with our interpretation of heterogeneity in marginal costs causing the bias in prices. Details of this robustness check are in the Appendix.

Third, another unique feature of the dataset is the fact that we observe exits by drivers. Each driver on Blablacar has a unique ID; thus, we can revisit profiles of drivers in order to check how active they have been after we have observed them. This allows us to establish whether drivers with high objective marginal costs are more likely to exit, and what is the impact of the reputation system on exit. Exploiting the exit data, we note that drivers who have less than four new ratings in December 2018 compared to the last ride we observe in our dataset (up to August 2018) set significantly higher prices than those who have more than 4 new ratings (that is posted at least one more time since their last observed ride).

The three arguments presented above support a claim that the differences in costs both marginal and opportnity are key factors in understanding the changes in the composition of the population of drivers, as well as their individual pricing decisions.

3.5 Results of the structural model

In this section, we first present the results of the estimation of the demand model. Second, we present the simulation of the supply side.

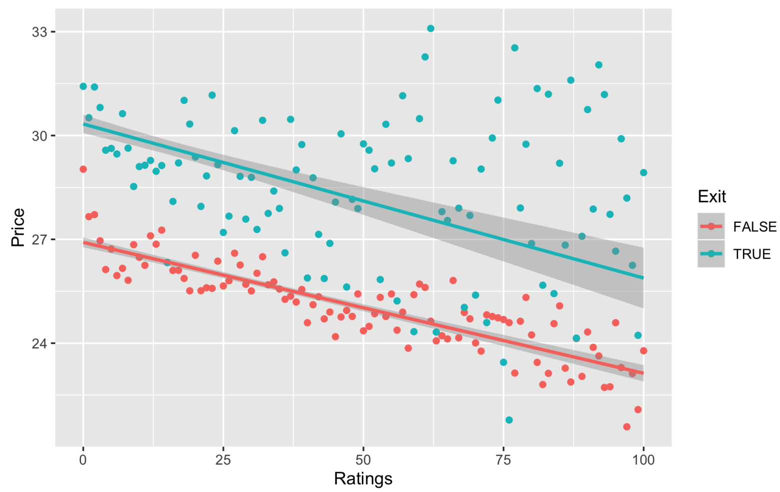

Figure 3.4: Prices of drivers who had posted at least 1 new listing by January 2019 since their last observed listing (in pink) are consistently lower than the prices of drivers who did not post a new listing since their last ride (in blue).

3.5.1 Demand side

Potential passengers search for rides on a route of their interest and observe available drivers. To receive detailed information about a driver, and to proceed to a booking stage, a passenger has to click on the listing. For every ride in our dataset, we observe this number of clicks/ views. We use this information as a proxy for the market size. In this section, the market size $m_{r,t}$ is assumed to be equal to the maximum number of views any listing in a route-day has generated. The estimation is conducted on 10 routes including all available days. The estimates in the table below show that passengers are price-sensitive with an average price elasticity of -1.38.

Using the estimated demand model and the observed market sizes, we recover sellers' marginal costs assuming full information per-period Nash equilibrium in prices, and in a first pass, static pricing. Prices set by sellers on the platform are per seat, so if the per-market Bertrand-Nash assumption holds, we can recover the marginal cost (per passenger) of a driver from the FOC of

Table 3.3: Conditional logit demand estimates

	Dependent variable:
	choice
Price	-0.05^{***} (0.004)
Ratings (log)	0.32^{**} (0.14)
Listings (log)	-0.14^{***} (0.02)
Average Rating	0.02 (0.08)
Ratings (log) x Average Rating	-0.01 (0.03)
Observations	2,594,650
Wald Test	$228{,}872.20^{***}$ (df = 49)
Note:	*p<0.1; **p<0.05; ***p<0.01

Within driver	0.61
Across drivers within route	3.62
Across drivers Across route	13.89

Table 3.4: Variance decomposition

the seller problem:

$$\hat{c}_{j,r,t} = p_{j,rt} + \frac{D_{j,r,t}}{\partial D_{j,r,t} \partial p_{j,r,t}} = c_r + \xi_j + w_{j,r,t} \tag{3.3}$$

We then use the estimated marginal costs (per passenger) to derive a simple test of the validity of the seller conduct assumption. The intuition of the test is simple: under Bertrand-Nash with full-information (no learning) and myopic agents, the shock to the marginal cost of seller is uncorrelated to either the number of listings a seller has posted, or the number of reviews they have received up to period t: $E(w_{j,r,t}n_{j,r,t}) = 0$ and $E(w_{j,r,t}nl_{j,r,t}) = 0$. Findings are in Table 3.5. The estimated cost shocks exhibit some correlation with the number of trips listed by the driver, an indication about the presence of learning.

In order to substantiate our previous claims on the selection taking place on the platform, we offer some descriptive results on the correlation between marginal costs and the probability of exit. The estimated marginal costs as shown in the table below are a strong predictor of the probability of exit of a

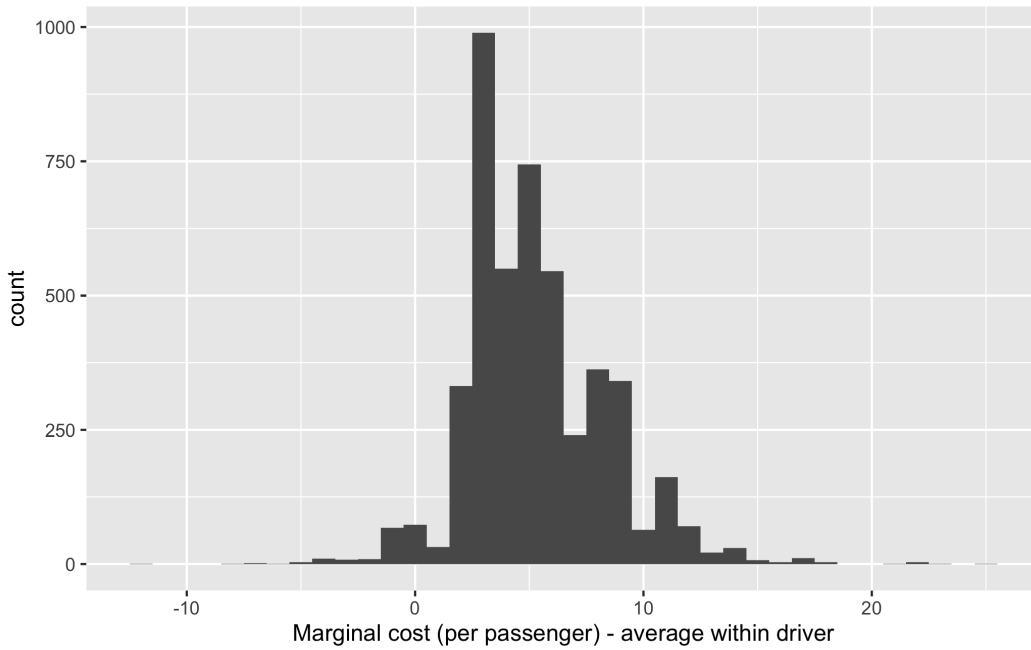

Figure 3.5: Costs of drivers on Paris-Angers

Table 3.5: Marginal cost regressions

	Dependent variable:	
	Marginal cost (per passenger)	
	Pooling	Within
	(1)	(2)
Ratings (log)	−0.504*** (0.059)	0.190 (0.124)
Listings (log)	−0.170*** (0.065)	0.517*** (0.162)
Average Rating	0.080*** (0.011)	−0.019 (0.020)
Route effects	X	X
Weekend, SNCF strike	X	X
Driver unobservable FE		X
Observations	9,555	9,555
Adjusted R^2	0.951	0.758
Note:		*p<0.1; **p<0.05; ***p<0.01

driver. A driver with a higher marginal cost is less likely to make a new listing on the platform even after controlling for the probability of sale at that price.

Table 3.6: Probability of exit regressed on the marginal cost

	Dependent variable:					
	Exit					
	(1)	(2)	(3)	(4)	(5)	(6)
Marginal cost (per passenger) c_j		0.041^{***} (0.007)		0.041^{***} (0.007)	0.022^{***} (0.007)	0.022^{***} (0.007)
No sales	0.158^{***} (0.057)	0.119^{**} (0.058)			0.008 (0.062)	
Revenue			-0.004^{***} (0.001)	-0.004^{***} (0.001)		-0.001 (0.001)
Average rating	0.194^{*} (0.105)	0.202^{*} (0.105)	0.199^{*} (0.105)	0.202^{*} (0.105)	0.004 (0.101)	0.004 (0.101)
Ratings (log)					-0.667^{***} (0.022)	-0.666^{***} (0.022)
Route effects	X	X	X	X	X	
Observations	9,546	9,442	9,468	9,442	9,442	9,442
Log Likelihood	−4,955.883	−4,878.708	−4,904.485	−4,876.714	−4,349.097	−4,348.812
Akaike Inf. Crit.	9,935.767	9,783.416	9,832.971	9,779.427	8,726.195	8,725.624
Note:						*p<0.1; **p<0.05; ***p<0.01

3.5.2 Simulations of the supply model

In this section, we simulate solutions to the theoretical model. For the sake of simplicity, we introduce several assumptions that make the model easy to solve; however, this still allows us to study how costs determine entry, exit, and pricing decisions. We are particularly interested in the role played by the opportunity costs? How much heterogeneity in opportunity cost is needed to generate the exit patterns observed in market studies (for example, JP Morgan Chase & Co Institute (2018)), or in the reduced form findings.

We hold fixed a competitive setting : 3 sellers compete in a market. The demand follows a logit specification with elasticities of price and reviews. The game is dynamic in the sense that sellers build a reputation from period to period, and we allow for exit (exit decision are driven by comparing all expected future profits with the opportunity cost). Reputation follows a probabilistic transition rule; *ex-ante* sellers have the same quality, and their expected first grade is the same, the past realizations of the grades influence the future grades. Thus, we assume that sellers have an intrinsic quality that is revealed in time, it is, however initially unknown both to the seller and to the market.

Model with opportunity costs

First, we simulate solutions to the model with opportunity costs. Timing of the game is as follows:

1. Sellers observe their reputation, the reputation of their competitors, and **opportunity costs**, which are drawn at random each period. They decide whether to stay at the platform or exit. They exit when the discounted sum of future profits is lower than the opportunity costs. We do not allow for re-entry.

2. Sellers set prices and profits are realized.

3. Reputations of sellers who stayed are updated.

We performed 600 simulations of Nash solutions to the demand problem and exit decisions. Sellers set Bertrand prices and exit the market if the draw of opportunity cost is higher then the value of continuation. In this sense, the exit decisions incorporate the static-dynamic tradeoff, while pricing is myopic. Figure 3.6 show results.

Figure 3.6: Simulated price paths

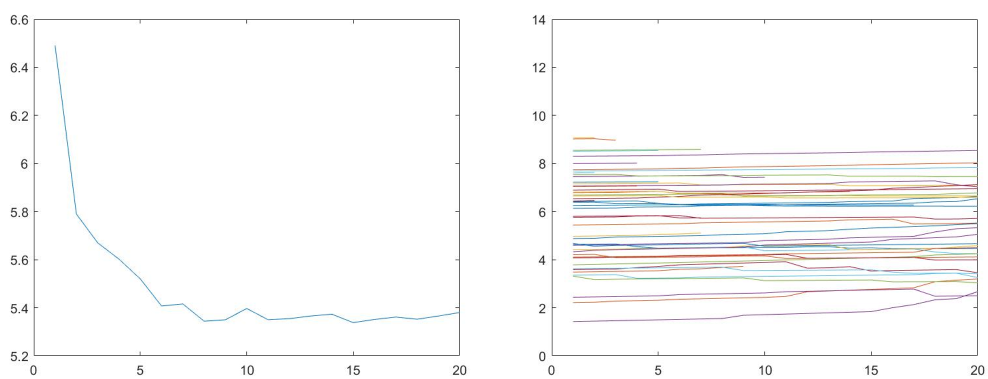

Note: The left panel average price in the market. The right panel shows an example of evolution of prices: randomly selected markets

Figure 3.6 left panel shows the average market price. We see that the average price is initially high and gradually decreases. Close inspection of the right panel reveals what influences such behavior. Initially, sellers with a high marginal cost set higher prices; however, they leave the platform the moment they receive a high draw of the opportunity cost. Sellers with low marginal costs remain on the platform and start building a reputation, which allows them to start gradually increasing the prices.

In order to uncover the effect of reputation building on prices, we focus on sellers who stayed for at least 20 periods. Figures 3.7 and 3.7 show the average prices by drivers who stayed pricing paths of some randomly selected drivers.

Figure 3.7: Simulated price paths incumbents

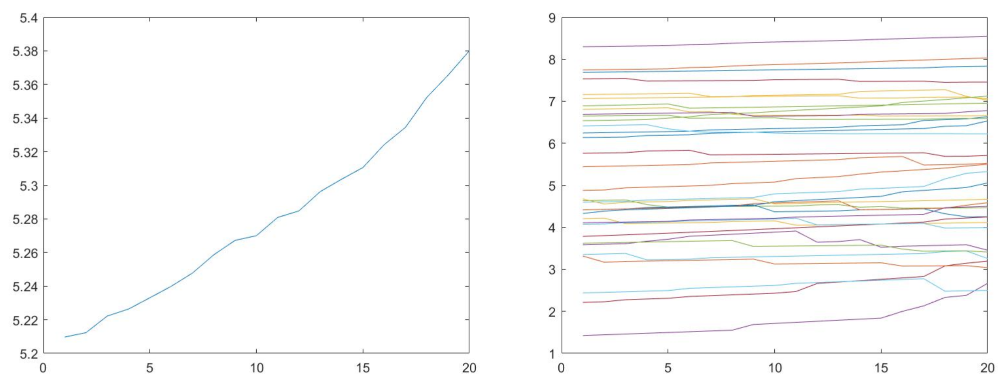

Note: The left panel average price of sellers who stayed for 20 periods. The right panel shows an example of evolution of prices of sellers who stayed

We draw two noteworthy conclusions from this simulation exercise: heterogeneity in marginal costs is crucial in generating price dispersion. Combination of opportunity cost and marginal costs shape exit patterns. Prices are initially high due to sellers with high marginal costs and decrease as they exit the market. Conditioned on staying on the market, a myopic entrant had set a lower price in the early periods than in the following periods - once she has build reputation, the optimal price is higher. Finally, observing changes in the distribution of prices and exit patterns, we can deduce the distribution of marginal costs and draws of opportunity costs.

Model without heterogeneity in opportunity costs

In this subsection, we introduce a change to the baseline model: we remove heterogeneity in opportunity cost. This exercise should allow us to understand better the role played by opportunity cost. Sellers differ in marginal costs. The timing of a period is the following:

1. Sellers observe their reputation, the reputation of their competitors, **shocks to their marginal cost** which are drawn at random each period. They decide whether to stay at the platform or exit. They exit when the discounted sum of future profits is lower than the opportunity costs (the same for all sellers and constant over time). We do not allow for re-entry.

2. Sellers set prices and profits are realized.

3. Reputations of sellers who stayed are updated.

We simulate per period Nash prices, as set by myopic agents, and exit decision is a cut-off rule. Figures below show the results of 600 simulations.

Figure 3.8: Simulated price paths

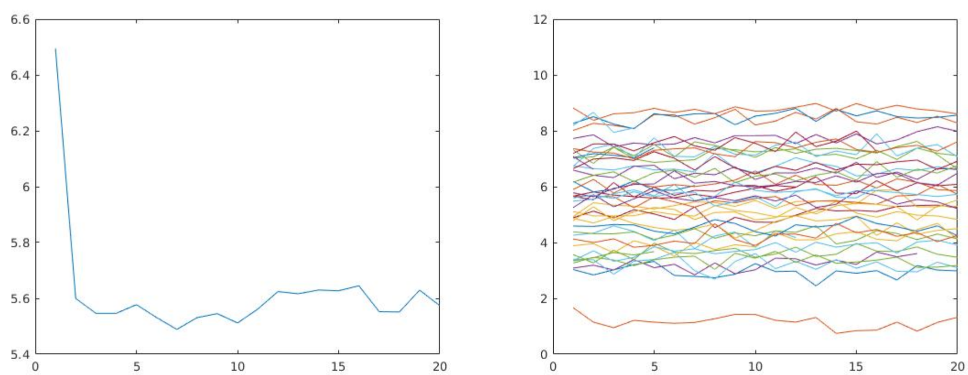

Note: The left panel average price. The right panel shows an example of evolution of prices.

Figures 3.8 and 3.8 show evolution of market prices (all sellers). We see that similar patterns emerge despite no heterogeneity in opprotunity cost. Prices are initially high and decrease gradually as inefficient sellers exit the market.

Figure 3.9: Simulated price paths

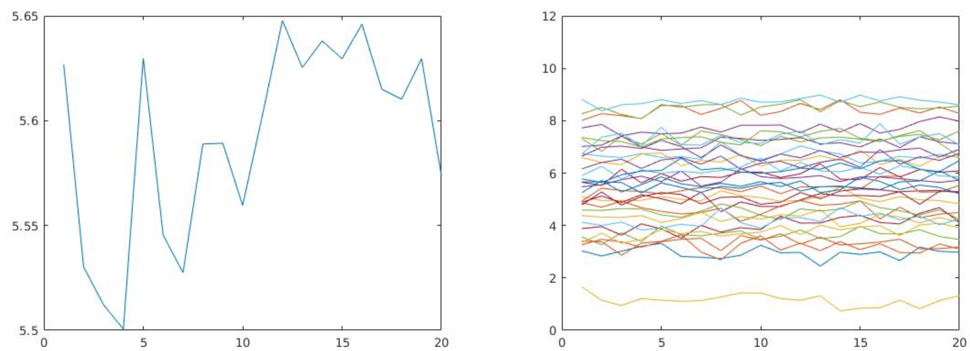

Note: The left panel average price of sellers who stayed for 20 periods. The right panel shows an example of evolution of prices of sellers who stayed

As in the previous simulations, we focus now on sellers who stayed for at least 20 periods. We observe the same patterns despite no difference in opportunity costs across sellers.

The second set of simulations reveals an interesting observation. Individual heterogeneity in opportunity costs is essential to understand individual entry and exit dynamics, however, average price dynamics are shaped by the difference in marginal costs. Sellers with the high marginal cost set higher prices and are more likely to exit. Efficient sellers start with low introductory prices and increase them gradually as they receive reviews. These patterns emerge without dynamic pricing by sellers. Arguably, in a fully dynamic model, efficient sellers with including a further price reduction to speed up the reputation building process, which in equilibrium will result in the faster exit of inefficient sellers. We are currently working on incorporating this dynamic in our model.

3.6 Conclusion

In this paper, we propose a model to study entry, exit, and pricing decision of sellers in a marketplace with a reputation system. We show that the heterogeneity in marginal and opportunity costs is the key to understanding the evolution of prices, as well as exit decisions. We, first, provide simulations of our model and, second, show that reduced form analysis of pricing on a popular ride-sharing platform has to account for the unobserved heterogeneity, to uncover expected price patterns.

Sellers with high opportunity cost set on average higher prices, and are more likely to leave the platform. Sellers with lower opportunity costs initially set low prices, and gradually increase them when they have built a reputation.

We are currently working on extending this draft in two dimensions. First, as mentioned before, a unique feature of our empirical set-up is that we can revisit profiles of drivers and measure whether they are still active or decided to stop using the platform. We want to extend our reduced form results by a study of exit decisions. Initial results reinforce our previous findings: sellers with high marginal costs (high travel costs) are more likely to leave the platform; also, a negative review has an effect of increasing the probability of exit. Second, our supply-side model rests on a number of assumptions that do not match the reality of a ride-sharing platform well. We are working on extending the model and our simulations to allow for a large and changing number of competitors and their strategic entry decisions.

Finally, during the period at study, Blablacar introduced a policy of promoting entrants by granting higher prominence to their listings. The effects of such a policy depend on how sellers make their pricing decisions and the distribution of unobserved costs. Our structural model allows us to estimate the effects of this policy change.

Bibliography

Bergemann, D. and Välimäki, J. (1996). Learning and Strategic Pricing. *Econometrica*, 64(5):1125–1149.

Bergemann, D. and Välimäki, J. (2000). Experimentation in Markets. *Review of Economic Studies*, 67(2):213–234.

Bergemann, D. and Välimäki, J. (2006). Dynamic Pricing of New Experience Goods. *Journal of Political Economy*, 114(4):713–743.

Besanko, D., Doraszelski, U., and Kryukov, Y. (2014). The economics of predation: What drives pricing when there is learning-by-doing? *American Economic Review*, 104(3):868–97.

Besanko, D., Doraszelski, U., Kryukov, Y., and Satterthwaite, M. (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, 78(2):453–508.

Cabral, L. and Hortacsu, A. (2010). The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics*, 58(1):54–78.

Doraszelski, U. and Pakes, A. (2007). A framework for applied dynamic analysis in io. *Handbook of industrial organization*, 3:1887–1966.

Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of economic studies*, 62(1):53–82.

Fershtman, C. and Pakes, A. (2012). Dynamic games with asymmetric information: A framework for empirical work. *The Quarterly Journal of Economics*, 127(4):1611–1661.

Hagiu, A. and Wright, J. (2018). Platforms and the exploration of new products. *Management Science (forthcoming)*.

Jolivet, G., Jullien, B., and Postel-Vinay, F. (2016). Reputation and prices on the e-market: Evidence from a major french platform. *International Journal of Industrial Organization*, 45:59–75.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica: Journal of the Econometric Society*, pages 649–670.

JP Morgan Chase & Co Institute (2018). The online platform economy in 2018: Drivers, workers,sellers, and lessors. Technical report, JP Morgan Chase & Co Institute.

Lambin, X. and Palikot, E. (2020). The impact of online reputation on ethnic discrimination.

Livingston, J. A. (2005). How valuable is a good reputation? a sample selection model of internet auctions. *Review of Economics and Statistics*, 87(3):453–465.

Maskin, E. and Tirole, J. (1988a). A theory of dynamic oligopoly, i: Overview and quantity competition with large fixed costs. *Econometrica: Journal of the Econometric Society*, pages 549–569.

Maskin, E. and Tirole, J. (1988b). A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica: Journal of the Econometric Society*, pages 571–599.

Miguel Villas-Boas, J. (2006). Dynamic Competition with experience goods. *Journal of Economics & Management Strategy*, 15(1):37–66.

Nelson, P. (1970). Information and Consumer Behavior. *Journal of Political Economy*, 78(2):311–329.

Pakes, A. and McGuire, P. (1992). Computing markov perfect nash equilibria: Numerical implications of a dynamic differentiated product model.

Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *the RAND Journal of Economics*, 38(2):373–399.

Spagnolo, G. (2012). Reputation, competition, and entry in procurement. *International Journal of Industrial Organization*, 30(3):291–296.

Weintraub, G. Y., Benkard, C. L., and Van Roy, B. (2008). Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411.

Weintraub, G. Y., Benkard, C. L., and Van Roy, B. (2010). Computational methods for oblivious equilibrium. *Operations research*, 58(4-part-2):1247–1265.

3.7 Appendix

Supply side estimation strategy

The key methodological challenge is to estimate the value of entering the platform from observed market outcomes. A seller takes into account the impact of entering the market today on the likely update to her state, but the value of that depends on future states of the market as well as her future draws of the opportunity cost. In this section, we propose a simplification of this problem based on a claim that if the dataset is rich enough it encompasses all recurrent industry states: in our case, this translates into profits of a seller x_j under the different realization of (X, z), and the observed transition probabilities $p(x'|x)$. The estimation strategy that we suggest has two steps:

1. Compute averages of the entry values (i.e., per period profits for sellers) and transition values $p(x'|x)$, which is simply the probability of obtaining a good or a bad review in period $t+1$ conditioned on a number of good or bad reviews in period t. These together allow us to calculate expected profits conditioned on being active in future periods.

2. We take the estimated values from Step 1 and treat them as actual expectation of sellers, to recover the distribution of opportunity costs ϕ. For a given market we observe several realizations of (z, X), which together with our estimates of $VC(x_j, X, z)$ allows us to characterize the probability density function of opportunity cost $f(\phi)$.

We can rewrite $VC(x, X, z, \phi; \theta)$ as

$$\mathbf{E}_{\phi',z',x'}\left[\max\left\{\phi' + \delta\mathbf{E}\left[VC(z'', x')|z', x'\right], \mathbf{E}\left[\pi(z', x')|z, x\right] + \delta\mathbf{E}\left[VC(z'', x'')|z', x'\right]\right\}\right]$$

,where a' denotes the realization of a random variable a in the next period, and a'' two periods ahead. The maximization problem can be represented as the choice probability problem, to obtain

$$
\begin{aligned}
VC(x, X, z, \phi; \theta) = Pr \{ &\phi \\
&< \mathbf{E} [\pi(z', x')|z, x] + \delta \Delta VC(z', x')\} \\
&\quad \cdot [\mathbf{E} [\pi(z', x')|z, x] + \delta \mathbf{E} [VC(z'', x'')|z', x']] + Pr \{\phi \\
&> \mathbf{E} [\pi(z', x')|z, x] + \delta \Delta VC(z', x')\} \cdot \mathbf{E} [\phi'|\phi' \\
&> \mathbf{E} [\pi(z', x')|z, x] + \delta \Delta VC(z', x') + \delta \mathbf{E} [VC(z'', x')|z', x]]
\end{aligned} \tag{3.4}
$$

,where $\Delta VC(z', x') = \mathbf{E} [[VC(z'', x'')|z'', x'] - [VC(z'', x')|z'', x']]$, is the expected gain in the value of being on the platform due to one review, the expectation is with respect to probabilities of obtaining a good or a bad review conditioned on being in state x'.

As an illustration, we assume that the opportunity costs follow exponential distribution $F(\phi) = 1 - e^{-1/\sigma\phi}$, which implies that $\mathbf{E} [\phi'|\phi' > A] = A + \sigma$. Thus, we can simplify equation 3.4 to write that:

$$
\begin{aligned}
VC(x, z, \phi; \theta) = \mathbf{E}_{z',x'} \Bigg[&\pi(z', x') + \underbrace{Pr [\phi'|\phi' > \mathbf{E} [\pi(z', x')|z, x]]}_{p(\phi')} \sigma \\
&+ \delta (VC(x', z', \phi; \theta) + (1 - p(\phi'))\Delta VC(z, x)) \Bigg]
\end{aligned} \tag{3.5}
$$

Denote by M the matrix of transition probabilities as perceived by the sellers: such that element (i, j) denotes probability of transition from a state (z_i, x_i) to a state (z_j, x_j), so it captures demand process $z_i \to z_j$, as well as reputation building: $x_i \to x_j$, which entails probability of getting a high review or a low one. We can represent the equation 3.5 in the matrix form:

$$
VC(\theta) = M [\pi + p(\phi')\sigma] + M [\delta (VC(\theta) + (1 - p(\phi'))\Delta VC(\theta))] \tag{3.6}
$$

We assume that, the number of good and bad reviews is bounded; we interpret this assumption that after a seller has collected enough reviews subsequent ones are not usefull for passengers. Thus, for each seller in state $l > L$, $\Delta VC(z, x_l) = 0$.

By plugging the defintion of $VC(\theta)$ into equation 3.6 and iterating we can get rid off the $VC(\theta)$ from the right hand side and obtain:

$$\begin{aligned}
VC(\theta) &= \sum_{\tau=1}^{L} \delta^{\tau} M^{\tau} \left[\pi + \sigma p(\phi') + \delta(1 - p(\phi'))\Delta VC(\theta)\right] + \sum_{\tau=L}^{\infty} \delta^{\tau} M^{\tau} \left[\pi + \sigma p(\phi')\right] \\
&= \sum_{\tau=1}^{\infty} \delta^{\tau} M^{\tau} \left[\pi + \sigma p(\phi') + \delta(1 - p(\phi'))\Delta VC(\theta)\right]
\end{aligned} \tag{3.7}$$

,where the last equality is obtained by noting that $\Delta VC(z, x_l)$ takes values of zero in a subset of states. Finally, the equation 3.7 can be further simplified to obtain:

$$VC(\theta) = \left[I - \delta M\right]^{-1} M \left[\pi + \sigma p(\phi') + \delta(1 - p(\phi'))\Delta VC(\theta)\right]$$

The first step in the estimation provides us with estimates of $\pi(z, x)$ for all (z, x), gain in profits $\Delta VC(\theta)$ is also estimated in the first step. Transition probabilities are directly observed in the data, the unknown parameter that has to be esimated is σ.

Definition and discussion of main variables

name of a variable	description
price	price set by the driver in EUR; has to be lower than maximum price: 0.082 per km
suggested price	price suggested by Blablacar: 0.065 per km
age	age of the driver in years
reviews	number of reviews received by the driver
male	gender defined based on photo recognition and name
minority	takes the value of one when the driver is of Arabic or African origin, and zero otherwise; defined based on photo recognition and name (see. Lambin& Palikot (2019)) for details)
picture	takes the value of one when driver added a picture, and zero otherwise
talkative	categorical variable (bla, blabla, blablabla) indicating how talkative the driver is
bio	number of words in driver's description
ride description	number of words in ride's description
reputation	mean of grades received by the driver
published rides	number of rides ever published by the driver
number of clicks	number of clicks a given listing has received; clicking is necessary for booking a ride but not sufficient; measured at the moment of data collection
sold seats	number of seats already sold; measured at the moment of data collection
revenue	sold seats multiplied by price
posts per month	mean number of listings posted by the driver since she joined the platform
seniority	number of months since the driver joined the platform
competition	number of listings available on the same day on the same route
median revenue	mean of median revenues in cities of departure and arrival; source: INSEE
public transport	travelling time by public transport on the route at listings' departure time; source: Google API
train strike	SNCF official strike implicating a given route
value of car	price of a comparable car model in thousands of EUR; when a model of a car is not available mean price of a brand; source: ebay (scrapped data)
fuel consumption	mean fuel consumption of a model of a car; when model of a car is not available mean consumption of a brand; source: ADEME
length (km)	distance in km between cities of departure and arrival; souce: Google API
lengh (hours)	estimated driving time by a car on a given route and time; source: Google API
hours until departure	number of hours between data collection and a ride departure
posted since	number of hours between the posting of the listing and data collection
automatic acceptance	takes the value of one if booking requests are automatically accepted and zero if the driver chose to accept/reject requests manually
to fuel price	average price of a litre of diesel in a city of arrival in cents
from fuel price	average price of a litre of diesel in a city of departure in cents
toll viamich	total toll costs on a given route in EUR; source: https://www.viamichelin.com/
travel costs	mean of fuel costs multiplied by fuel consumption plus toll fees
weekday	takes a value of 1 on weekdays and zero on weekends
pets	takes a value of 1 if the driver accepts pets and zero otherwise
music	takes a value of 1 if the driver listens to music in the car and zero otherwise
smoke	takes a value of 1 if the driver accepts smoking in the car and zero otherwise
detour	categorical variable: 1 if no detour, 2 if some detour (up to 15 min), and 3 if more than 15 minutes detour
luggage	categorical variable: 1 if no luggage, 2 if small bags, 3 if big bags are allowes

Table 3.7: Definition of main variables

Table 3.8: Summary statistics

Statistic	N	Mean	St. Dev.	Min	Max
Price	211,164	26.979	12.933	5	80
Number of ratings	211,164	49.212	100.644	0	1,579
Number of listings	211,164	70.005	146.262	1	2,895
Average rating	211,164	4.595	0.275	1.000	5.000
Seats offered	211,164	2.649	0.767	1	4
Seats taken	211,164	0.327	0.634	0	4
Weekend	211,164	0.325	0.468	0	1
SNCF strike	211,164	0.042	0.201	0	1
Automatic acceptance	211,164	0.462	0.499	0	1

Table 3.9: Summary statistics- drivers

Statistic	N	Mean	St. Dev.	Min	Max
Last observed number of ratings	51,057	41.417	64.006	0	1,573
Last observed number of listings	51,057	54.460	89.588	1	2,104
Last observed average rating	51,057	4.591	0.253	1.000	5.000
Driver observations	51,057	2.104	3.027	1	117
Photo	51,057	0.883	0.321	0	1
Age	51,057	37.610	13.553	18	102
Male	51,057	0.688	0.463	0	1
Exit (0 new trips)	51,057	0.267	0.442	0	1
New trips	51,057	9.526	18.272	0	633

Table 3.10: Summary statistics - route/day

Statistic	N	Mean	St. Dev.	Min	Max
Number of listings	27,866	7.578	11.807	1	203
Mean price	27,866	30.036	14.325	5.000	80.000
SD price	20,557	2.991	2.393	0.000	42.426
Market size	27,866	48.772	58.381	0	1,398
Seats offered	27,866	20.073	31.272	1	565
Seats taken	27,866	2.476	5.255	0	133
Listings w/ 0 booked seats (share)	27,866	0.814	0.251	0.000	1.000
Listings w/ less than 5 reviews (share)	27,866	0.359	0.348	0.000	1.000
Listings w/ 1st trip (share)	27,866	0.130	0.266	0	1

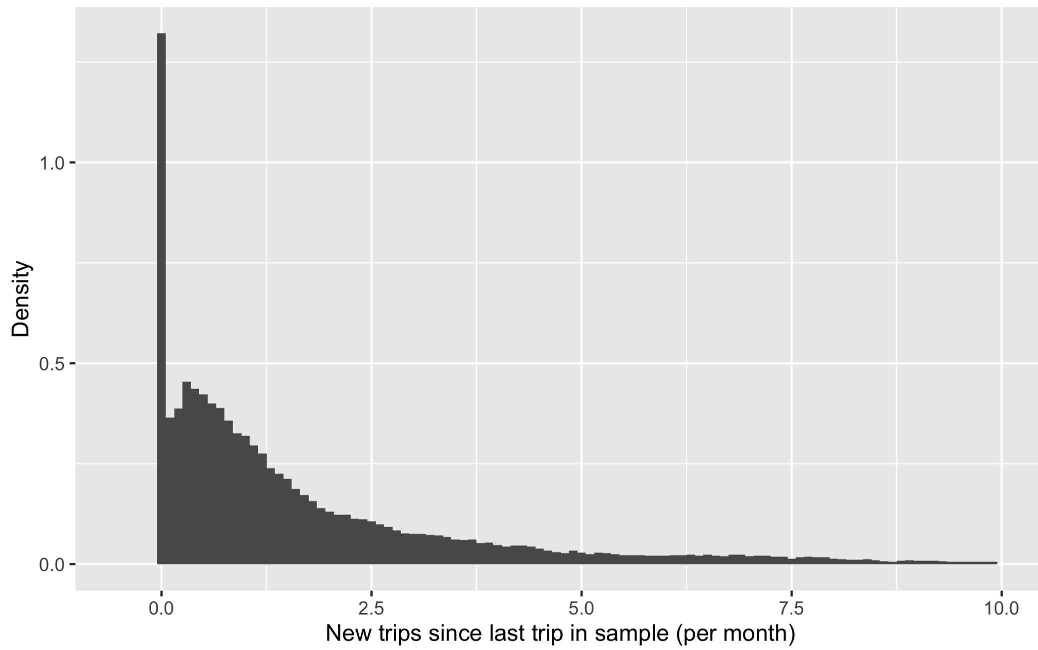

Figure 3.10: Distribution of number of new trips since last scraped listing (normalized per month).

Robustness checks

Robustness check for the sign of the bias: we restrict our dataset to drivers whom we had seen when they entered the platform (five reviews or less) and stayed until they collected at least 40 reviews. This sub-sample is smaller than the unrestricted one; thus we estimate the effect of a coarser measure of being an entrant: having 5 or fewer ratings. The results are reported in Table 3.11. The within estimates of the pricing behavior of entrants (in column 4) shows a is a negative and significant sign, although there is a loss in precision. However, what is interesting is that in the subsample of drivers whom we observe on different stages of their careers, the pooling estimator of entrant status is also negative and significant. In this subsample, sellers price lower in early posts than later on, even without controlling for unobserved heterogeneity. Drivers in this subsample stayed on the platform for a long period; thus the selection of the subsample is not random: it is conditional on the driver not exiting before having at least 40 reviews. This observation indicates that conditioning on not exiting, the correlation between the unobserved fixed effect and the number of ratings becomes close to zero. We formally test this by a Hausman test of the within estimator versus the random estimator. We cannot reject the null hypothesis that both estimators are consistent in this subsample, whereas we reject it at 1% risk level in the larger sample. The unobserved driver-specific fixed effects that are driving our bias in Table 3.1 are then highly correlated with exit behavior. This gives credence to our interpretation of these effects as the average opportunity cost of the driver of using the platform: drivers with high opportunity costs leave the platform early on in their career but are still selected in the repeated observation sample of Table 3.2, thus the difference between the random and within estimators. The subsample we select for Table 3.11 amounts to conditioning on non-exit, this means selecting the drivers with low average opportunity costs. In this subsample, the random effects estimator is very close then to the fixed effects estimator: these are drivers with low opportunity costs who start at a lower price and increase gradually.

Table 3.11: Sub-sample with drivers we observe early and late stage of career

	Dependent variable:			
	Price of a ride			
	OLS		*panel*	
	Pooling	Within	Pooling	Within
	(1)	(2)	(3)	(4)
Constant	56.841***		−6.328***	
	(1.188)		(2.143)	
Entrant	0.588***	−0.427***	−0.711***	−0.209*
	(0.019)	(0.043)	(0.232)	(0.118)
Average rating	−0.275***	0.009	1.973***	−0.226
	(0.030)	(0.085)	(0.312)	(0.341)
Travel cost	0.030***	0.057***	0.417***	0.052***
	(0.001)	(0.003)	(0.002)	(0.008)
Sample	Full	Repeat observations	Early-late	Early-late
Travel cost	X	X	X	X
Time effects	X	X	X	X
Driver Controls	X	X	X	X
Trip effects	X	X	X	X
Observations	238,397	168,520	9,342	8,712
R^2	0.930	0.891	0.802	0.917
Adjusted R^2	0.930	0.842	0.802	0.905
Residual Std. Error	4.122 (df = 238111)			
Note:				*p<0.1; **p<0.05; ***p<0.01