ABSTRACT

Title of dissertation: REPUTATION DYNAMICS
IN MARKETING CONTEXTS

Kalinda Ukanwa Zeiger
Candidate for Doctor of Philosophy, 2019

Dissertation directed by: Dr. David Godes
Professor and Department Chair
Marketing Department

Dr. Roland Rust
Distinguished University Professor and
David Bruce Smith Chair in Marketing
Marketing Department

This dissertation is an examination of the impact of dynamic consumer reputation effects on firm decision-making in the marketplace.

Essay I is a study of the impact of firm interventions on competitive reputation building among consumers on an online platform. Specifically, I model an actor's decision to upload pirated content in order to build his reputation, despite facing threats from copyright lawsuits (firm interventions seeking to deter uploading activity) and intense competition on the platform. We propose a novel theory that explains what could occur in this scenario: high-reputation consumers will decrease their reputation-building activity, but their low-reputation competition may see an opportunity to enhance their reputation and increase activity. We argue that because competition for reputation is active on the site, the lawsuits may deter

uploading in the short-run but may actually lead to *more* piracy over the long-run. Our findings support the theory: while high-reputation publishers decrease the likelihood of uploading as lawsuits increase, low-reputation uploaders do the opposite: they upload more.

Essay II is an examination of the impact that a consumer group's reputation can have on firm decision-making. Specifically, we investigate 1) conditions under which a non-prejudiced firm may discriminate in service against its consumers based on group reputation, and 2) how subsequent consumer word-of-mouth can impact demand and profits over time. This mixed-methods study shows that discrimination can be profitable in the short-run but can backfire in the long-run due to the effects of consumer word-of-mouth and firm competition. Results indicate that high consumer heterogeneity in quality (i.e., their profitability to the firm) and low measurement error in detecting consumer quality attenuate the magnitude of service discrimination. The authors provide managerial recommendations on reducing service discrimination's profit-damaging effects. This research emphasizes the long-term benefits of switching to a service policy that does not use group reputation information.

This dissertation contributes to the general marketing literature by providing new insights into how the reputation of the consumer, a sparsely researched area, can have direct impact on the firm in its decision-making.

REPUTATION DYNAMICS IN MARKETING CONTEXTS

by

Kalinda Ukanwa Zeiger

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2019

Advisory Committee:
Professor David Godes, Chair
Professor Roland Rust, Co-Chair
Professor John Chao
Professor P.K. Kannan
Professor Michel Wedel
Assistant Professor Bo (Bobby) Zhou

Acknowledgments

I am grateful to the many individuals who made this dissertation possible. First and foremost I would like to thank my advisor, Dr. David Godes, for giving me an invaluable opportunity to work on extremely interesting projects during the program. His invaluable advice, mentorship, and all around support has been critical to my growth as a researcher. It has been a pleasure to work with and learn from such a great scholar and extraordinary individual.

I would also like to thank my co-advisor, Dr. Roland Rust. His guidance, sponsorship, and persistent drive for excellence enabled me to build and stretch scholarship muscles that I did not know I had. It has been a game-changing experience working with Dr. Rust, who has been a fantastic ally and mentor.

Thanks are also due to Dr. P.K. Kannan, Dr. Michel Wedel, Dr. Bo "Bobby" Zhou, and Dr. John Chao for agreeing to serve on my dissertation committee. Using their invaluable time to share their expertise and provide feedback on this manuscript has significantly boosted the quality of this research.

I want to express my appreciation to the faculty, students, and staff in the Marketing Department at the University of Maryland for all their helpful advice, feedback, and support. Special thanks to the students in my writing groups for their weekly feedback and counsel, including Jared Watson, In Hye Kang, Min Kim, and Xian Gu.

Most importantly, many, many thanks to my husband Lenny and my daughter Ariana. Their support, encouragement, patience, and incredible personal sacrifices were the foundation for me to pursue and complete this dissertation. I could not have done it without them.

Finally, I would like to acknowledge financial support from the Marketing Science Institute and from the office of the Dean of Smith School of Business for the projects discussed herein.

Table of Contents

List of Tables

List of Figures

Chapter 1: Competition for Reputation

1.1 Introduction

> *The power of the reputation effect depends on the nature of one's opponents; notably on whether they also seek to acquire a reputation. – Kreps and Wilson (1982, p. 275)*

Prior literature has shown that reputation can be a valuable resource or asset (Klewes and Wreschniok, 2009; Tadelis, 2002, 2003). A number of studies have established the benefits of possessing a strong reputation, such as gaining a competitive advantage (Fombrun and Shanley, 1990; Kreps and Wilson, 1982; Milgrom and Roberts, 1982), increased value for one's products, services, or offerings (Bajari and Hortaçsu, 2004; Resnick et al., 2006; Yoganarasimhan, 2013), increased demand (Cabral and Hortaçsu, 2010; Chevalier and Mayzlin, 2006), advantage in acquiring more reputation (Merton, 1968, 1988), and enhanced trust from others (Resnick and Zeckhauser, 2002). Other studies have shown the prospect of building a reputation is a significant motivator of consumer participation in online marketplace activities (Dellarocas et al., 2004; Diekmann et al., 2014; Taylor, 2012b).

However, what is not well understood are 1) conditions under which consumers compete with each in the building of their marketplace reputations and 2) how firm interventions (attempts to change consumer behavior) can impact the reputation-building dynamics. Our research investigates the decision to acquire reputation in a competitive context. Furthermore, we examine how firm interventions can impact the competitive dynamics on a platform. Organizations often use interventions to deter consumers from targeted behaviors. A substantial literature on incentives and deterrence across disciplines would predict that the threat of a penalty will decrease targeted behavior (Bandura, 1969; Ehrlich, 1996; Schwartz, 1989). However, other studies have demonstrated that deterrence interventions can increase targeted behavior (Carroll et al., 2009; Gneezy and Rustichini, 2000; Mendoza et al., 2017; Taylor, 2012a). We propose a new theory, driven by an unexplored contextual setting: that competition for reputation (a social dynamic), can differentially influence individual response to firm interventions and produce more of the targeted behavior. To explore this phenomenon, we model the reputation-building decisions of whether to upload unauthorized content to an online piracy platform. The uploaders, known as publishers, face decision tradeoffs due to two countervailing forces: competition from other publishers and reputation-building costs due to threats posed by thousands of copyright lawsuits seeking to deter online piracy (see Appendix B for more on how lawsuit campaigns generate awareness).

The top graph in Figure 1.1 displays weekly uploads during the period of January 1, 2010 to December 31, 2013 on Kickass Torrents, the online platform that is the focus of our study. The bottom graph displays BitTorrent-related U.S. Federal copyright lawsuits filed during the same period. Kickass Torrents was at the time one of the most popular digital file sharing sites in the world. Note from these

Figure 1.1: Growth in Site Uploads vs. BitTorrent Copyright Lawsuits

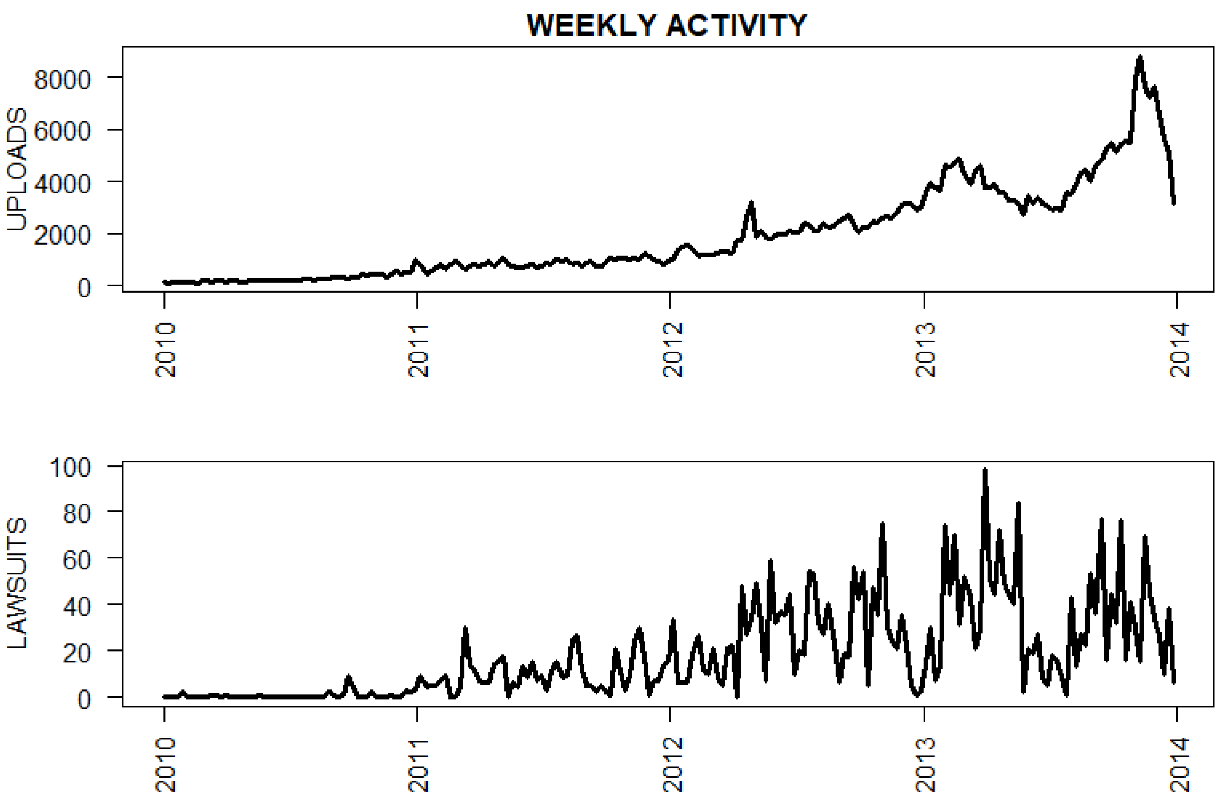

graphs that despite the increasing volume of lawsuits during this period, there is a concomitant increase in uploads to the Kickass Torrents site. It is not immediately apparent from the graphs whether the lawsuits are effective at reducing piracy.

Qualitative evidence suggests that some copyright owners or trade groups behind these lawsuits used them strategically as a means of deterrence by raising awareness of the illegality of piracy and the penalty of doing so. Indeed, the Recording Industry Association of America (RIAA) had in 2013 the following quote on their site under a section called "WHY WE DO WHAT WE DO":

> Prior to the lawsuits, only 35 percent of people knew file-sharing was illegal, but after the initiation of the end-user legal campaign, that number quickly rocketed to more than 70 percent. ...If awareness of the copyright laws and an appreciation of the consequences of getting caught for breaking the law had not had an effect, p2p growth rates would likely have continued unabated, and would have seriously undermined the potential for a legal digital marketplace. (RIAA, 2013)

Furthermore, the lawyer for Malibu Media LLC, which filed 38% of all copyright lawsuits in the US in 2015, stated:

> From our experience, if you file a certain number of lawsuits in a particular district, there's no question that the infringement goes down. ... The more articles there are in the press, the more aware that people are that what they're doing can expose them to liability, the more deterrence there is. This kind of deterrence requires a large volume of lawsuits, [the lawyer] said. You can't sue one or two people occasionally and expect that to have an effect on the infringement. (Mazumdar and Shen, 2016)

To examine how lawsuits as a deterrence intervention impact competition for reputation among pirating individuals, we model and analyze the decisions of hundreds of participants who upload unauthorized content to the platform during this period. We find evidence consistent with our proposed theory that competitive reputation-building may impact the decision to engage in the targeted behavior (uploading in this case). Deeper analysis reveals that the lawsuits appear to be more effective with high-reputation uploaders, who decrease uploading activity as the intensity of lawsuits increases. However, low-reputation uploaders appear to increase their activity in response. One implication is that managers and policy-makers should take consumer competitive social interactions into consideration when designing firm interventions meant to discourage specific behaviors.

This study contributes new knowledge about consumer competitive behavior as a function of reputation and its implications for firm interventions. We elaborate on these themes in the remainder of the paper as follows: in Section 1.2, we discuss prior literature and our contribution. We present our conceptual framework in Section 1.3. In Section 1.4, we discuss our empirical context and the data, followed by explanation of the empirical model in Section 1.6. Analyses and results are discussed in Section 1.8. We discuss research limitations and conclude in Section 1.9.

1.2 Related Literature on Reputation, Deterrence, and Piracy

This research contributes to literatures about reputation, firm interventions (specifically, deterrence measures), and online piracy. Prior theoretical work on reputation demonstrates that a positive reputation can be a valuable asset or resource (Klewes and Wreschniok, 2009; Tadelis, 2002, 2003) that creates cumulative advantage stratification over time between those with high and low reputation (Cole et al., 1974; Merton, 1968, 1988). The theoretical literature has also established conditions under which firms are willing to make costly investments to build a reputation Fudenberg and Levine (1989, 1992); Kreps and Wilson (1982); Milgrom and Roberts (1982). However, firms with higher levels of reputation find less benefit in investing than those with low levels Board and Meyer-ter Vehn (2013). A comparatively smaller body of work provides empirical support for these assertions (Cabral and Hortaçsu, 2010; Jin and Leslie, 2009; McDevitt, 2011). This prior literature on reputation focuses on the reputation of firms. In contrast, our work examines the reputations of consumers and their implications for the firm or marketplace. Furthermore, there is little in the literature that accounts for competitive dynamics in reputation building. Two exceptions are Mayzlin et al. (2014) and Luca and Zervas (2016), who show that firms lacking a strong reputation are more likely to create fake online reviews in an effort to bolster their own reputations relative to their competition's. We differentiate ourselves from these studies in that we provide insight into how competition for reputation itself can change incentives to acquire reputation. Furthermore, we provide insight into conditions where costs that are heterogeneous in reputation levels can change the dynamics of reputation building. These important differences are gaps that our study seeks to close.

There is a large literature on deterrence and disincentives, particularly in the areas of law, economics, and information systems security. Deterrence theory presumes that to deter individuals from taking an action, a firm or organization must make the cost of taking the action sufficiently high enough and the likelihood of capture credible enough to outweigh the benefit of acting (Becker, 1968; Ehrlich, 1996; Stigler, 1970). For example, threats of expulsion should deter club members from breaking rules, high fines should deter customers from paying their bills late, and threats of lawsuits should deter consumers from engaging in piracy. Similar logic is proposed in the threat and fear appeal literatures, which proposes that the threat imposed raises a salience of cost which could exceed the benefit of taking an action and thereby reduce the action (Dickinson-Delaporte and Holmes, 2011; Maddux and Rogers, 1983; Rogers, 1975; Tanner et al., 1991). Much of the prior literature regarding deterrence considers the decision-making of the targeted individual without accounting for social considerations (Becker, 1968; Ehrlich, 1996; Stigler, 1970). Other literature that accounts for social considerations examines individual response to deterrence measures as a function of peer pressure (Carroll et al., 2009). However, there is a small and growing literature that shows that deterrence measures can produce unintended consequences. For example, prior studies have shown deterrence interventions backfiring due to economic motivations, which thereby increased the number of tardy parents picking up their kids from a day-care center (Gneezy and Rustichini, 2000) and the incidents of tax evasion (Mendoza et al., 2017). Other studies propose psychological and sociological mechanisms for deterrence measures backfiring. For example, deterrence measures that increased incidents of smoking among teen-agers (Carroll et al., 2009) and increased frequency of graffiti activity (Taylor, 2012a) are attributed to reactance and peer conformity motives. Our study significantly differs from these explanations by proposing a new mechanism driven by the concept of competition between consumers to build reputation. We propose that competitive dynamics and the quest for reputation-building benefits can diminish the effects of firm interventions.

We also contribute to the literature on the effectiveness of anti-piracy legal interventions. Notable papers include Danaher et al. (2014), who find that increased consumer awareness of HADOPI, France's graduated response law, increased French iTunes music sales by 22-25%. Danaher and Smith (2014) investigate U.S. government copyright policy enforcement and find that law enforcement's shutdown of Megaupload.com, a major piracy site, led to a 6.5% increase in sales for the movie industry. Our paper is perhaps closest to Bhattacharjee et al. (2006), who find that the RIAA's strategy of announcing and then employing well-publicized legal actions to discourage music file sharing was only partially successful in decreasing music file sharing. While large file sharers substantially decreased their sharing volumes, small fire sharers decreased moderately, and very few completely stopped the activity all together. Furthermore, they found an upsurge in file sharing activity after the RIAA made its second announcement of lawsuits. Our study also finds that legal interventions (in the form of lawsuit campaigns) produce mixed results in reducing piracy behavior. However, our study differs from Bhattacharjee et al. (2006) in two important ways. Their study does not provide answers as to why different groups of file sharers respond differently to the legal interventions. In contrast, our study proposes and finds empirical support for a conceptual framework which explains why some file sharers react differently to publicized legal threats. We find evidence consistent with our proposed theory: that heterogeneous reputation incentives shaped by competition for reputation and costs drive the behavior. Second, their study does not address conditions where the legal threats can backfire and produce upsurges in activity. We find conditions consistent with our theory that can explain an upsurge in piracy activity. This not only has policy implications about effectiveness of legal campaigns but also has marketing implications for influencing volume of participation on online platforms. To the best of our knowledge, no prior study examines how an anti-piracy policy can change the dynamics of competition for reputation and unintentionally produce more piracy. We next turn to describing our conceptual framework in detail.

1.3 Conceptual Framework

Our conceptual framework concerns an agent's decision as to whether to take an action to build his reputation on a platform. Our framework applies to market contexts that meet four criteria: 1) the context provides a mechanism where reputation signals are made visible; 2) reputation is earned through actions taken by the agent; 3) there are other agents also seeking reputation; 4) there are positive payoffs to the agent as a function of reputation. Although our framework can apply to multiple settings that meet these criteria (e.g., e-commerce sites like eBay and Amazon, sharing economy sites like Airbnb), we ground our explanation of the framework in online piracy in order to match our empirical setting. The agent in this case is the publisher who wishes to develop a reputation for being a trustworthy supplier of unauthorized content on the piracy platform he participates in. Reputation has value because it engenders trust from other members in the community and could help attract downloads of his unauthorized content. To differentiate himself from hundreds of other publishers on the platform, the publisher also must compete with them by having higher levels of reputation. For example, imagine three publishers on an online piracy platform upload the same movie file on the same day. Without price (which is 0) or observable product quality differences to distinguish the publishers from each other, consumers on the platform may instead base their selection decision on the comparative reputation levels of the publishers. If consumers find that one publisher has 30% better reputation scores than the others, this may influence consumer likelihood to select content of the publisher with the reputation advantage.

In our context, we examine the reputations of publishers for uploading on the Kickass Torrent site where the observers are other publishers and downloaders of content. We define reputation as signals of an attribute which are a function of repeated, observable actions (Fombrun, 1995; Kreps and Wilson, 1982; Milgrom and Roberts, 1982). In our context, these signals are a function of the action of uploading that publishers make on the site. It is important to note that we do not examine quality indicators of the uploads, however (e.g., ratings, up votes vs. down votes, etc. of uploads). This treatment of a dimension of reputation as a function of actions that is orthogonal to the quality of the action is consistent with known reputation system design on online platforms, where reputation recognition can be given for actions taken on the site, quality of actions on the site, or a combination of the two (Farmer and Glass, 2010). Because we do not observe dynamic measures of quality of actions (e.g., ratings given per period or up votes given per period) on the site, we do not provide insights to competition for reputation for quality over time. We leave such an inquiry to future research.

Uploads are an appropriate measure of reputation for being a publisher because 1) it is a repeated action done over time that is observable by other publishers, downloaders, and by the researchers; 2) uploads have a strong relationship with publisher reputation points overall[1] on the site. After ranking publishers by cumulative uploads on the platform and cumulative reputation points, we found that the two rankings had a 0.68 correlation. Hence, we model publisher reputation as a function of uploads he makes on the platform.

However, reputation is a costly resource to build and has diminishing returns (Board and Meyer-ter Vehn, 2013; Resnick and Zeckhauser, 2002). The publisher faces a decision in each period of whether to invest in an action to acquire reputation on the platform. An example of an action could be uploading a copy of the latest "Star Wars" to the piracy site. Given that he wishes to maximize his reputation, the participant's key considerations in his decision include his current reputation level, his rivals' reputation levels, his future potential actions, and his competitors' future actions to acquire reputation. These factors determine the likely benefit he will derive from acquiring additional reputation. He weighs this against the cost to acquire new reputation.

We assume that marginal costs to acquire additional reputation are a function of current reputation levels, consistent with cumulative advantage theory (DiPrete

[1]We do not use reputation points in the analysis because we do not observe dynamic data on point accumulation. Furthermore, participants can earn reputation points for activity other than uploading Kickass Torrents. Our focus, however, is on the reputation of being a publisher who uploads content to the site.

and Eirich, 2006; Merton, 1968, 1988). For example, a publisher may experience decreasing costs to acquire new reputation as a function of his current reputation because his reputation affords him better access to resources for new content to upload. On the other hand, a publisher may experience increasing costs as a function of his reputation because his greater visibility may attract greater risks of getting caught by law enforcement or copyright owners. Therefore, the publisher (and agent in general) needs to consider how acquiring additional reputation will affect his acquisition costs.

The publisher faces a trade-off: he weighs the returns to taking an action to acquire new reputation as a function of the level of competition he faces in the same period versus the marginal cost to acquire it. If there are other publishers uploading higher volumes of content on a given day, his uploads may not get the same attention as on a day where there are fewer uploads. Less attention means his action has less impact on increasing his reputation. Furthermore, we assume that the publisher knows that acquiring new reputation today can change his future acquisition costs, since costs are a function of reputation level. He realizes that other competing publishers face the same trade-off about acquiring reputation. Thus, the publisher must form an expectation of what his rivals will do in order to complete his assessment of his own potential gains. If the benefits to acquire reputation outweigh the costs under these conditions, he takes an action to acquire new reputation in the current period.

1.4 Data and Statistics

Kickass Torrents, one of the sources of our data for our study, was launched in 2008 as a social media platform centered on BitTorrent file sharing. Before a joint U.S. Department of Justice/Homeland Security/IRS raid shut down its servers in July, 2016, Kickass Torrents was the most popular BitTorrent[2] site in the world (Van der Sar, Ernesto, 2016). Kickass Torrents allegedly facilitated the distribution of over $1 billion in unauthorized content and had upwards of 50 million visitors per month (Blake, Andrew, 2016).

Figure 1.2 is a snapshot from a page on Kickass Torrents' site. It provides qualitative evidence that its participants seek to build their reputations. At the top of this figure is a logo created by Yify, a popular publisher on Kickass Torrents. Yify also made sure to create visibility for himself by including his publisher name at the end of the file names of his uploads (three files on the left hand side, circled in red). Furthermore, the site creates a platform that emphasizes attributes signaling reputation. In this example, the site visibly displays the number of uploads, badges associated with uploaders, and designations of verification and level of uploader on the publisher Yify's profile page.

The data derived from Kickass Torrents is an unbalanced panel dataset of historical uploads from publishers listed on the site as "verified uploaders" at the

[2] BitTorrent is one of the most popular communication protocols in the world to distribute data and electronic files over the internet. It uses peer-to-peer file sharing of large data files, such as audio or video.

Figure 1.2: Uploader Reputation Example

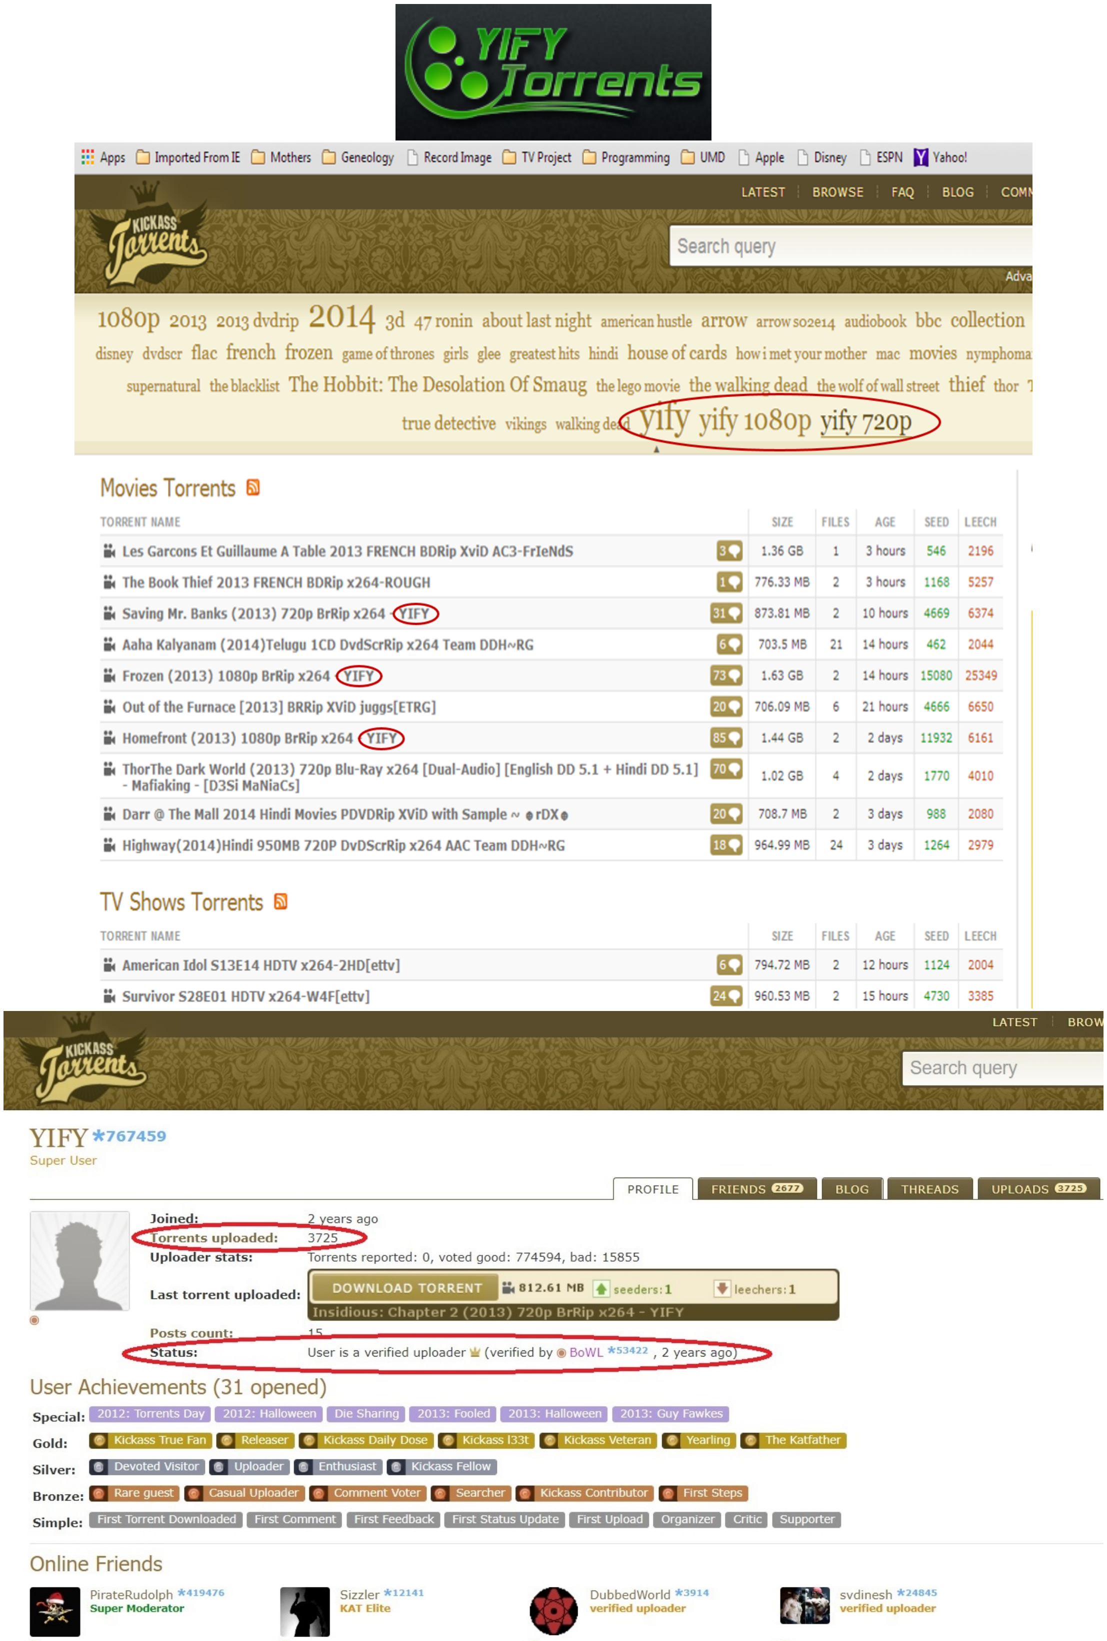

Source: Internet Archive Wayback Machine, Dec. 8, 2013

time the data was collected. "Verified uploaders" are site participants who had successfully passed a vetting process to upload in the community. Site administration allowed only "verified uploaders" to upload content to the site. We collected the data by scraping the log pages of uploads provided by individual publishers. Table 1.1 provides summary statistics of the data on publisher activity. This dataset

Table 1.1: Summary Statistics: Focal Publisher

	Mean	SD	2.5%	97.5%
Days Since First Upload	661.9	592.7	41.8	2383.4
Cumulative Sessions	137.4	168.5	3.0	645.5
Cumulative Uploads	533.9	1251.8	4.0	3623.0
Uploads/Session	4.4	12.3	1.0	21.6
Focal Publisher's Competitors/Session	102.1	52.2	14.0	186.0
Focal Publisher's Competing Uploads/Session	430.5	289.8	26.0	1148.0

NB: Session = day the publisher uploaded at least once.

includes 752 publishers who uploaded files tagged to one and only one of nine content categories: Anime, Books, Games, Movies, Music, Applications, TV, XXX, and Other (a category that usually contains posters, magazines, etc.). Each upload record includes a date-time stamp, name of the content uploaded, and content category. Transactions included in the data set occurred during the period of January 1, 2010 through December 31, 2013.

This represents a 4-year period of transactional daily data, which results in

1,461 days of transactions and over 450,000 records of published content. Given that each publisher's profile includes a historical record of every one of his uploads, this setting allows us to observe his activities across time. For this reason, this setting is particularly compelling to study the dynamics of reputation-building activities.

Publishers, on average, upload 84% of their content in a single content category, indicating that they tend to specialize in content category. Nearly half of the publishers (46%) in the dataset specialize in movies or music. The typical publisher has had almost 2 years of experience publishing on the site. He has an upload session, on average, every 4 - 5 days, in which he uploads four items per session. A session is defined as a day on which a publisher uploads content at least once.

The other primary source of data collected in this study were records of copyright lawsuits filed in U.S. Federal Court under U.S. Code 17:101 (the Copyright Act). We augment the Sag (2015) dataset with additional data sourced from the U.S. government database called Public Access to Court Electronic Records (PACER). There were over 1,800 copyright lawsuits filed during the period of the study. Each record in our data contains the court filing date, the state and federal district the suit was filed in, and whether the lawsuit violation occurred in BitTorrent. Since the identity of many of the defendants is indeterminate in the law suits (often defendants are listed as John and Jane Does), we do not observe whether the lawsuits specifically target any of the Kickass Torrents publishers in our dataset. However, we do observe evidence of publicity and communications about the lawsuits. We also observe evidence of word of mouth about lawsuits in the piracy community (see Appendix B for examples of this). Therefore, we assume the lawsuit campaigns raise general awareness of the consequences of piracy and that awareness has an impact similar to the effects of advertising (Chung, 2013; McCormick and Tinsley, 1987; Murphy and Trandel, 1994). Furthermore, we assume that lawsuits act as shocks to the cost of acquiring new reputation via uploading unauthorized content. In the next section, we present model free evidence related to our research goals of understanding how competition and firm interventions in the form of lawsuits may affect uploading activity to acquire reputation.

1.5 Model-Free Evidence

Figures 1.3 through 1.5 provide model-free evidence of the relationship between the decision to upload and the environmental conditions of competition and lawsuits. All graphs display data on a weekly basis during the period of the study. The x-axis of the first four graphs shows total uploads on the Kickass Torrents website during a given week. The reason why we examine total uploads is that it is a measure of the intensity of competition on the site during a given week. The greater the number of total uploads on the site, the smaller the share of voice (attention) for his uploads for any given publisher. The last two graphs show lawsuits filed in a given week. The y-axis on the first two graphs shows the number of sessions per week conducted by the average high-rep and low-rep publisher respectively. The last four graphs show the percentage change of average number of sessions per week for high-rep and low-rep publishers on the y-axis. Note that first graph suggests that as the total uploads on the site initially increases (competition increases), there is a corresponding sharp

Figure 1.3: Publisher Avg. Sessions per Week vs. Total Uploads to Site

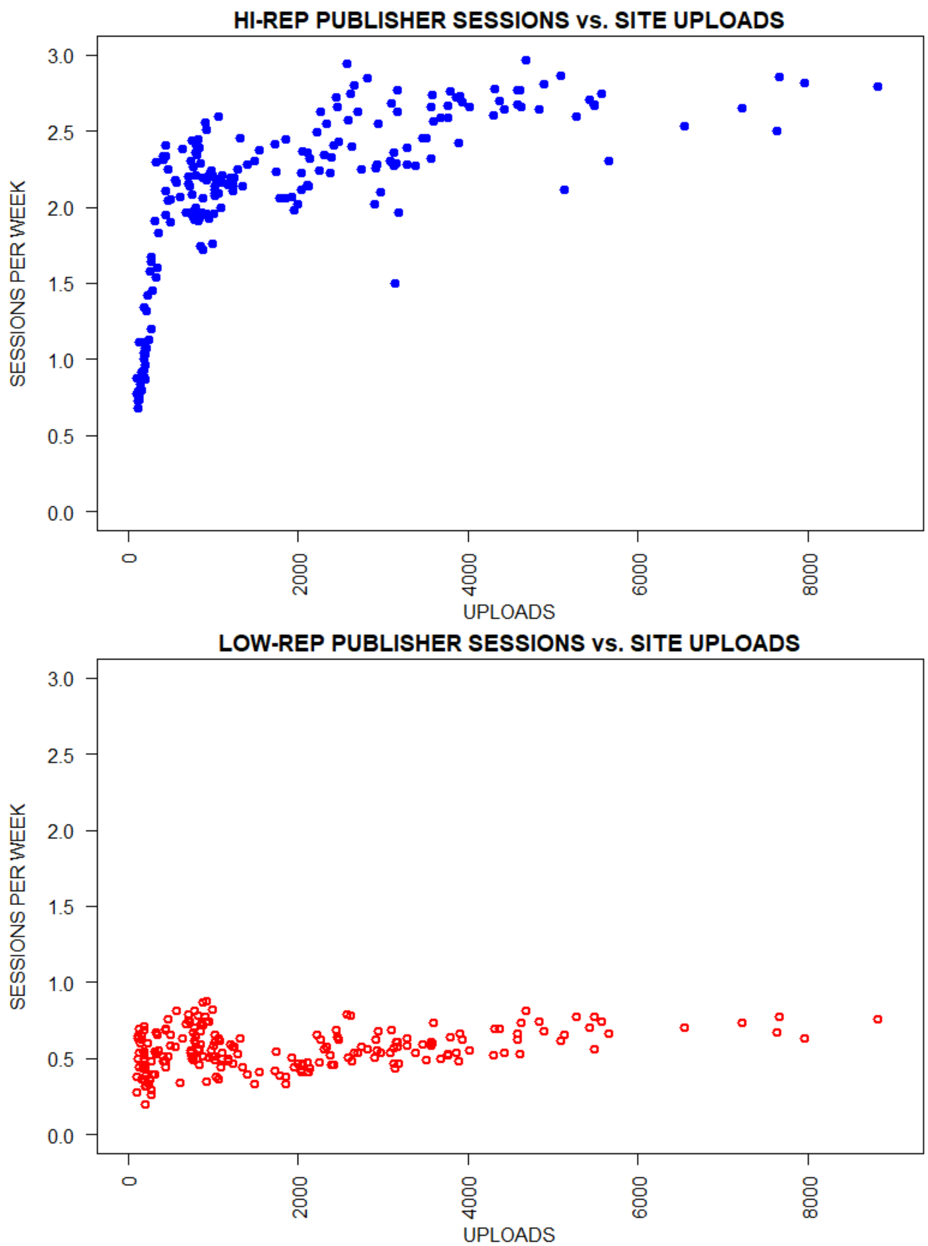

rate of increase of weekly sessions per high-rep publisher. However, as total uploads continue to increase past 1000, the rate of weekly sessions begins to decelerate sharply. In contrast, the second graph shows that low-rep publishers only slightly increase their weekly sessions as competition increases. Upon examination of how

Figure 1.4: Change in Avg. Sessions per Week vs. Total Uploads to Site

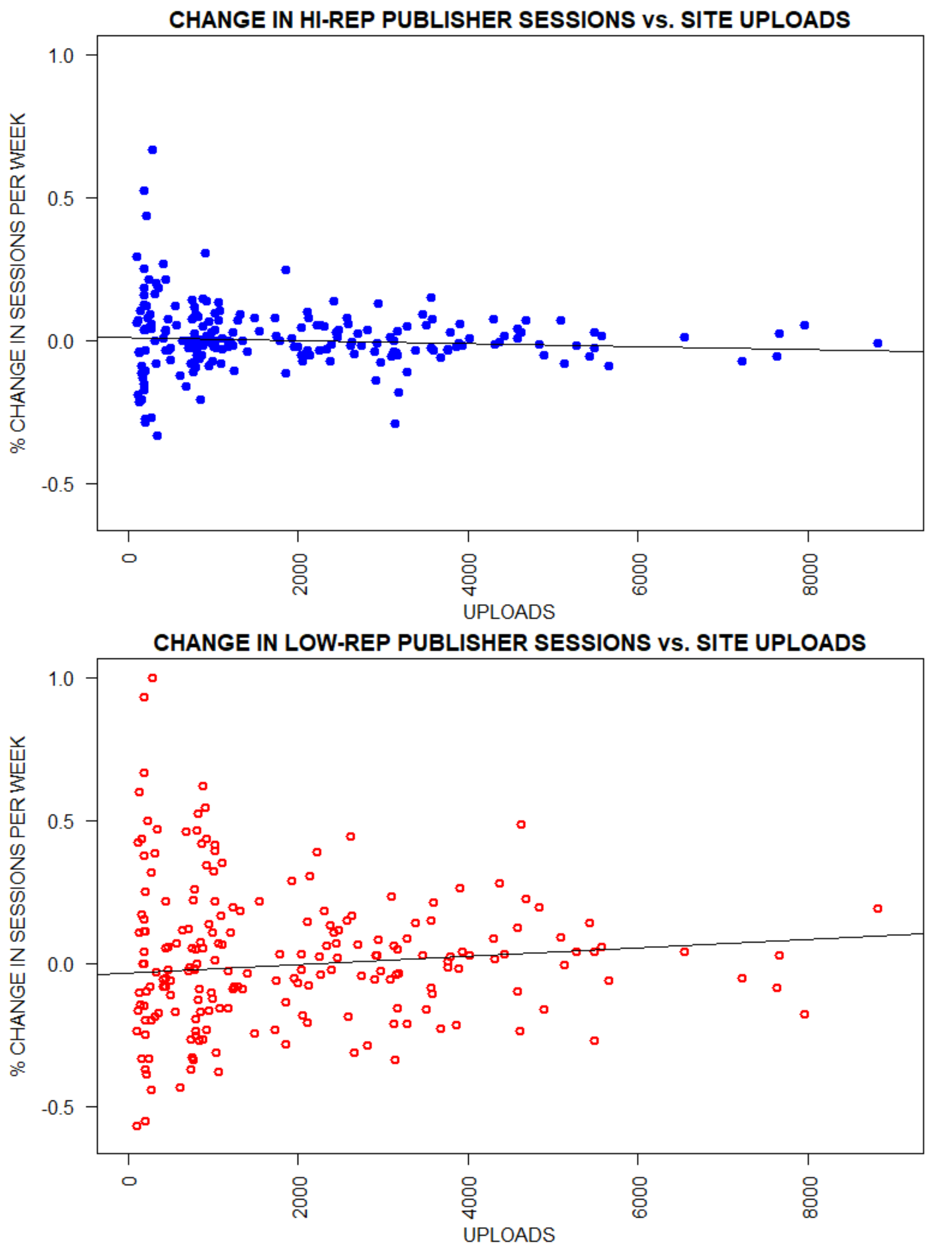

change in average sessions per publisher correlates with total uploads on the site

(Figure 1.4), we gain the additional insight that both high-rep and low-rep publishers have greater magnitude of changes when total uploads on the site are low. As total uploads increase, however, there is a corresponding dampening of the magnitude in changes of weekly uploading sessions. Figures 1.3 and 1.4 collectively suggest that increasing intensity of competition may be associated with a decrease in publisher

Figure 1.5: Change in Avg. Sessions per Week vs. Total Lawsuits Filed

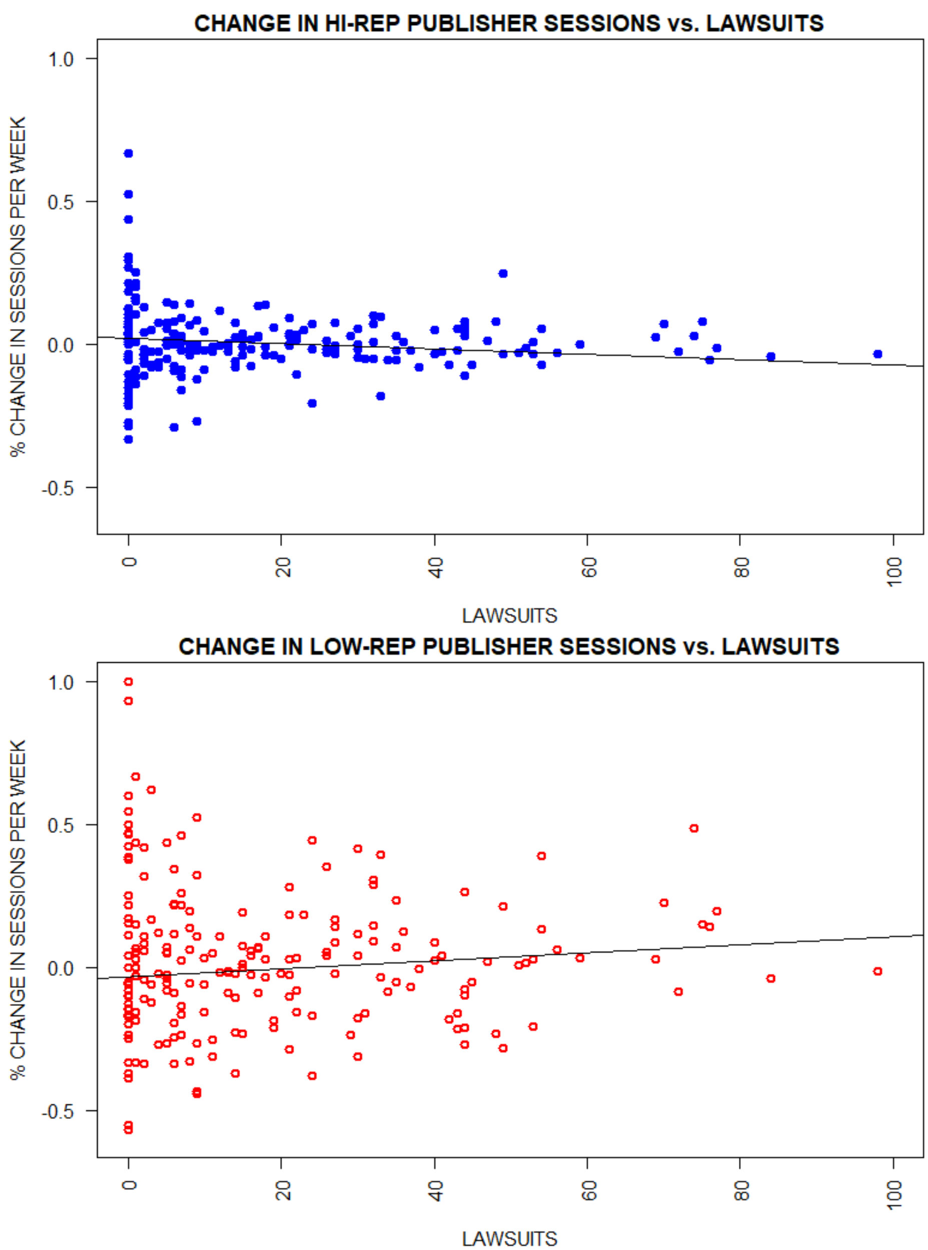

session activity. The graphs also suggest that such a decrease is more pronounced for hi-rep than low-rep publishers.

Figure 1.5 compares changes in session activity relative to increasing rates of lawsuit filings. The graphs in this figure suggest that as lawsuits increase, for high-rep publishers there is a corresponding decrease in percentage change in weekly sessions. On the other hand, there is a corresponding increase among low-rep publishers. The data indicate that lawsuits may have a heterogeneous effect on publishers, depending on their reputation level. To explore these themes, we model the daily decision of a publisher of whether or not to upload unauthorized content to Kickass Torrents in the presence of increasing volumes of BitTorrent-related U.S. copyright lawsuits. In the next section, we discuss the details of the model.

1.6 Empirical Model

1.6.1 Publisher's Decision

We begin by modeling the publisher's decision of whether to acquire reputation in the current period by uploading unauthorized content to the platform. Let us assume that publisher i seeks to maximize his utility (U_{it}) associated with building his reputation at time t. He will take this action–i.e, he will upload content–in order to build his reputation if the marginal utility to do so exceeds that from no action. The utility is comprised of the positive utility for increased reputation less the direct cost and expected penalty derived from uploading. The publisher's net utility for uploading unauthorized content in each time period is

(1.1)

$$U_{it} = Rep_{it}^{new}(Uploads_{it}) - c[Rep_{it}^{new}(Uploads_{it})]$$

$Rep_{it}^{new}(Uploads_{it})$ = New reputation acquired when i took action(s) $Uploads_{it}$ at time t

$c[Rep_{it}^{new}(Uploads_{it})]$ = i's cost to acquire new reputation at time t

We assume that the publisher's utility for reputation is proportional to the "share of voice" he expects to capture by uploading content. The more often he publishes content, the more awareness he generates.

We model reputation (Rep_{it}) as a stock variable renewed by actions ($Uploads_{it}$) to capture its enduring impact over time. This is because reputation, as a construct, is derived from past activities that can impact perceptions (Fombrun, 1995). The newly acquired incremental reputation enters a stock model which represents the reputation level of participant i. The stock model used is a commonly used functional form in the literature (Burmester et al., 2015; Dubé et al., 2005; Simon, 1982):

(1.2)

$$Rep_{it} = \ln(1 + Uploads_{it}) + \lambda^R R_{i,t-1}$$

λ^R = the time-invariant reputation carryover rate which is common to all i

Prior literature on reputation has shown that returns to reputation decline in reputation. Board and Meyer-ter Vehn (2013) show conditions where those with low and intermediate levels of reputation have incentives to invest or work to build higher levels of reputation, but those with high reputation or more likely to shirk due to declining incentives. Obloj and Obloj (2006) find evidence in a Polish auction site that the value of reputation diminishes in reputation levels, especially when the focal reputation level is sufficiently differentiated from the nearest competitor's reputation. Livingston (2005) and Yoganarasimhan (2013) find that value of reputational building decreases in additional ratings on an eBay and an online freelancer site (respectively). Other papers on Internet auctions show that bid premiums decrease in reputation levels, thus indicating a diminishing returns to reputation (Bajari and Hortaçsu, 2004; Eaton David H, 2007; Lucking-Reiley et al., 2007; Resnick and Zeckhauser, 2002; Resnick et al., 2006). In our context, we examine the utility of performing an action (uploading) to build reputation. Given substantial evidence through prior work that the value of reputation decreases in reputation level, we expect to find that utility for performing an action, conditional on already having a high-reputation level, will be less than for someone who has low reputation level. This would imply that the value for taking an action to build new reputation declines in reputation. Hence, we expect the coefficient on the reputation variable to be negative.

The utility of reputation can be impacted by competition from other publishers. The greater the volume of uploaded content from competing publishers at the time the focal publisher uploads content, the less share of voice the focal publisher earns. Thus, we model utility for building reputation as a function of the actions of the focal publisher as well as that of his competitors. Furthermore, current levels of reputation can impact the value of newly acquired reputation. Because their reputations are already well-established, highly reputed individuals often derive less utility for a reputation-building action than their counterparts with little to no reputation (Board and Meyer-ter Vehn, 2013). We model the payoff of newly acquired reputation as a function of taking action to acquire reputation, individual i's current reputation, and the reputation and actions of his competition.

The publisher's cost to acquire reputation has two major components: a direct cost to upload content and a cost associated with the risk of punishment. The direct cost to publish content includes any effort and financial investment required to find, acquire, and publish content on the site. The direct marginal cost to publish an additional unit of content may be constant or non-constant in reputation of the publisher. A case for constant costs would suggest that the same amount of effort and financial costs are needed for each post, regardless of the reputation of the publisher. An argument for costs that decrease in reputation (cumulative advantage in costs) would suggest that search, acquisition, and publishing costs could decrease in reputation because the publisher's high level of reputation can provide easier access to resources to acquire content. For example, if a publisher's high reputation results in other users voluntarily bringing content to the renowned publisher, then the publisher's cost to upload content has just decreased, thus decreasing his marginal reputation acquisition cost. Because these costs are a direct function of uploads, however, there could be a problem with identification of this cost. For this reason, we assume the cost is captured in the utility associated with reputation.

On the other hand, costs associated with the risk of punishment are arguably increasing in publisher reputation (cumulative disadvantage). These costs are a result of the risk of getting caught publishing unauthorized content. Publishers who are caught doing so face potential financial penalties from their Internet Service

Provider, the US law enforcement, and legal action by the copyright owner. Since being caught in the act of publishing unauthorized content is not a certainty, the publisher faces an expected punishment cost: the probability of capture multiplied by the expected penalty if captured. Penalties are assessed per item of content published to the site, so the total expected cost of punishment increases in unauthorized content items published. The penalty, if punished, can be quite high. Per U.S. Code 17:504, copyright owners can elect to settle with an infringer before final court judgment and collect between $750 to $30,000 per infringed upon copyrighted work (a single act of uploading copyrighted work is an infringement). After a court judgment in favor of the copyright owners, that amount can go up to as high as $150,000 per infringement (U.S. Copyright Office, 2018). Furthermore, the marginal cost of expected punishment could also be increasing in reputation. This is because the publisher with high levels of reputation may attract more retaliatory attention by rivals and stakeholders such as law enforcement and copyright owners, thereby increasing costs (Weber, 2009). Each additional publication, which drives incremental reputation, increases the attention risk.

We assume that lawsuits are shocks to the cost to acquire reputation. Given that we model an individual publisher's decision to acquire reputation, we assume lawsuits are exogenous; it is highly improbable that copyright owners and trade organizations decide to file BitTorrent copyright lawsuits because of an individual publisher's decision to upload. Because members of the piracy community often learn about the lawsuits through news media, copyright owner or trade association announcements, Internet Service Provider messages, and word-of-mouth (e.g., piracy site forums), we assume the impact is analogous to advertising affects. We model the impact of these lawsuits as a stock variable (Law_{it}) in a manner similar to Chung (2013) who argues that non-traditional source of awareness, such as college athletics, could produce impact akin to advertising effects.

$$\begin{aligned} Law_{it} &= \lambda^L Law_{i,t-1} + \ln(1 + l_t), \text{ where} \\ \lambda^L &= \text{a time-invariant lawstock carryover rate common to all } i \\ l_t &= \text{lawsuits filed at time } t \end{aligned} \tag{1.3}$$

However, unlike traditional advertising which seeks to motivate the targeted audience to do something, the effects associated with lawsuits intend to discourage the targeted audience. Similar to threat appeals, their intended effect is to raise the salience of costs associated with the reputation-building activities of publishing unauthorized content (Dickinson-Delaporte and Holmes, 2011; Maddux and Rogers, 1983). Thus, we expect the coefficient on the lawsuit stock variable to be negative.

We assume that the logit transformation of carryover rates follows a multivariate normal distribution.

$$\mathbf{logit}\lambda \sim \mathcal{N}(\mathbf{logit}\lambda_{\mathbf{0}}, \mathbf{\Sigma}) \tag{1.4}$$

$$\text{where } \mathbf{logit}\lambda = \begin{pmatrix} logit(\lambda^R) \\ logit(\lambda^L) \end{pmatrix}$$

By modeling the distribution of carryover rates jointly, we allow for possible correlation between sensitivity to information about lawsuits and sensitivity to the impact of his reputation. For example, if publishers are more sensitive to the lasting impact of their reputations (possesses a high reputation carryover rate), they may also be more vigilant about protecting their reputations. Therefore, they may actively seek information about BitTorrent lawsuits and be more sensitive to exposures to such information (possesses high lawstock carryover rate). By modeling this potential correlation between reputation effects and lawstock effects, we can then control for it in our understanding of the effects of reputation stock and lawstock on reputation acquiring behavior (Braun and Moe, 2013).

The interactions of reputation, lawstock, and expected share (competition) are the variables of interest that enter the publisher's utility for acquiring reputation. The interactions are important in answering the following questions raised by our conceptual framework: 1) do the effects of lawsuits (lawstock) impact publishers asymmetrically with respect their reputation levels? 2) does competition differentially impact publishers with respect to their reputation levels? 3) how does the publisher's decision change when he considers his competition with respect to lawstock levels?

We use rational expectations theory (Lucas, 1976; Muth, 1961) as a theoretical basis to model publisher i's expectation of the number of uploads competitors will publish in the current period, $E(CompUploads)_{-it}$. This implies the following assumptions: 1) that all publishers form expectations of future competitive uploads based on past and present information available to all on the platform; 2) on average, publishers' expectations do not systematically differ from realized platform outcomes. In other words, we assume that $E(CompUploads)_{-it} = CompUploads_{-it} + \epsilon$, where ϵ is i.i.d. and $E(\epsilon) = 0$.

We assume that the publisher expects his competitors have similar considerations (volume of expected competitive upload activity and intensity of lawsuit awareness). The publisher's expectation of competitor's uploads is a function of the prior period's total uploads from competitors ($CompUploads_{-i,t-1}$), lawsuits ($Lawsuits_{i,t-1}$), and their interaction. We include weekly, day of week, month, and year fixed effects are also taken into consideration to control for unobserved shocks that are common to all publishers (e.g., major holidays, technical issues on Kickass' website, etc). We also include Gaussian copula regressors ($EndogCompUpload_{-i,t-1}$, $EndogLawsuits_{i,t-1}$) to control for potential endogeneity stemming from the CompUploads and Lawsuits variables (see Section 1.6.2 for more details on endogeneity and Park and Gupta (2012) Gaussian copulas). We use the following dynamic Tobit model to represent the publisher's expectation of the number of competing uploads:

(1.5)

$$\begin{aligned} E(CompUploads)_{-it} = \beta_0 &+ \beta_1 CompUploads_{-i,t-1} + \beta_2 Lawsuits_{i,t-1} \\ &+ \beta_3 CompUploads_{-i,t-1} \cdot Lawsuits_{i,t-1} + EndogCompUpload_{-i,t-1} \\ &+ EndogLawsuits_{i,t-1} + \text{time-based fixed effects} + \epsilon \\ \text{where } \epsilon &\sim \text{Tobit Type I} \end{aligned}$$

We chose a Tobit model, as opposed to alternative models such as the Poisson or Negative Binomial regression model because it is more consistent with our theoretical framework concerning the publisher's forecast of his competition in the current period. The Poisson and Negative Binomial regression models would not be appropriate for our study because they both model Poisson processes. Poisson processes inherently assume that the occurrence of events are independent of each other and of time. In our case, the event (competing uploads) is not independent of the time nor other's uploads. In fact, in our model specification, we assume that the prior period's competitive uploads and number of lawsuits directly inform the publisher's expectation of current competing uploads. Furthermore, the publisher's expectation is specifically taking into account the likely action of others. Hence, the Poisson and Negative Binomial assumption of independence between events is inherently violated.

In contrast, the dynamic Tobit model (Wooldridge, 2005) allows for the outcome to be at least partially dependent on previous outcome history. Furthermore, within our theoretical framework we assume that publisher's believe that the number of competing uploads will equal to 0 unless the volume of competitive uploads and lawsuits were not too high in the prior period. The magnitude of the volume of either in the prior period will directly influence the magnitude of the expectation in the current period. This is consistent with the theoretical framework in the usage of Tobit models in other studies such as labor supply or purchase of goods (Amemiya, 1973). Given this theoretical framework and statistical properties that support our context, the dynamic Tobit model is well suited to model the publisher's expectation of competitor's uploads.

We assume competition's impact on an uploader's utility function is in terms of share of voice: the reputational benefit of one upload is proportional to the fraction it comprises of total uploads posted to the site in that period. The measure of i's expectation of one upload's share of voice is represented as $EShare_{it}$:

$$EShare_{it} = \frac{100}{1 + E(CompUploads)_{-it}} \tag{1.6}$$

We use the inverse of expected competitor's uploads and scale it by 100 to facilitate estimation by making the covariate size comparable with other variables in the model. We use an inverse functional form because we assume that competition has a non-linear impact that declines in amount of competition. For example, if publisher i faces two other competitors in a given day and anticipates an additional competitor will upload that day, the impact on the publisher's share of voice in the community will be much greater than if i faces 200 competitors and anticipates an additional competitor will upload that day. Modeling competition in terms of share as a function of the inverse of the number of rivals (or in this case, competing uploads) is a common specification in the competition literature (Berry and Reiss, 2007). A positive coefficient on $EShare_{it}$ indicates that the publisher experiences greater utility from uploading when there is less competition. A negative coefficient on $EShare_{it}$ indicates that the publisher experiences greater utility from uploading when there is more competition, which may suggest a herding behavior where publishers like to publish when others do as well.

The final portion of our model specification includes daily fixed effects to control for unobserved daily shocks that are common to all publishers and Gaussian copula regressors ($EndogRep_{it}, EndogLaw_{it}, EndogShare_{it}$) to address endogeneity concerns associated with reputation, lawstock, and expected share (discussed in more detail in Section 1.6.2).

Publisher i's complete utility function associated with an upload is as follows:

(1.7)

$$U_{it} = \alpha_{i0} + \alpha_{i1} Rep_{i,t-1} + \alpha_{i2} Law_{i,t-1} + \alpha_{i3} EShare_{it} + \alpha_{i4} Law_{i,t-1} \cdot EShare_{it}$$
$$+ \alpha_{i5} Rep_{i,t-1} \cdot EShare_{it} + \alpha_{i6} Law_{i,t-1} \cdot Rep_{i,t-1} + EndogRep_{it} + EndogLaw_{it}$$
$$+ EndogEShare_{it} + \text{daily fixed effects} + \varepsilon_{it}$$

We use a Hierarchical Bayesian probit to model and estimate publisher i's decision. Modeling the decision with random coefficients will provide more insight into how publishers that are heterogeneous in reputation levels form their decisions conditional on competition, anticipated reputation-acquisition costs, and cost shocks imposed by the lawsuits. The dependent variable in the model is whether publisher i has an upload session (i.e., he uploads at least one item in time period t).

For the sake of compactness, we use matrix notation to represent the Hierarchical Bayesian probit model:

(1.8)

$$Pr(I = 1|X, D, T) = Pr(U > 0) = \Phi(U)$$
$$U_i = X_i \beta_{1_i} + D_i \beta_{2_i} + T_i \beta_\tau + \varepsilon_i$$
$$U = XB_1 + DB_2 + TB_\tau + E$$
$$\varepsilon_{it} \sim \mathcal{N}(0, 1)$$

$X =$ stacked X_i. The design matrix of publisher variables driving utility

$D =$ stacked D_i. Regressors based on Gaussian copulas to address endogeneity

$T =$ stacked T_i. Daily fixed effects common to all

$B_1, B_2, B_\tau =$ coefficients for X, D, and T respectively

$$
\begin{aligned}
B &= [B_1|B_2] \\
B &= Z\Theta + \Delta \\
\beta_i &= \Theta^{`} z_i + \delta_i \\
\delta_i &\sim \mathcal{N}_P(0, \Lambda)
\end{aligned} \tag{1.9}
$$

Z = a matrix of mean-centered time-invariant variables associated with publisher attributes.

Θ = matrix of coefficients on Z, representing average effects of average publisher

Δ = matrix of error terms

To complete the model specification, we employ diffuse priors[3] on the hyper-parameters of the Hierarchical Probit:

$$
\begin{aligned}
vec(\Theta'|\Lambda) &\sim \mathcal{N}_{PQ}(M_0, \Lambda \otimes A^{-1}) \\
\Lambda &\sim \mathcal{IW}_P(\nu_0, V_0) \\
A^{-1} &= 100 I_Q \\
M_0 &= 0_{PQ} \\
\nu_0 &= P + 3 \\
V_0 &= \nu_0 I_P
\end{aligned} \tag{1.10}
$$

P = number of variables in utility function

Q = number of variables in Z

[3]We also used flat priors to check robustness in this assumption and find that the results are qualitatively the same. Thank you to Dr. John Chao for this suggestion.

Conditional on the uploader's expectation of competing uploads in the current period, the resulting model's likelihood function for the uploader's decision to upload is as follows:

$$\begin{aligned}\mathcal{L}(y_i|\beta_{\mathbf{1_i}}, \beta_{2_i}, \beta_\tau, \lambda^R, \lambda^L) = \mathcal{N}(\Theta^{`} z_i, \Delta)\ \mathcal{N}_{PQ}(M_0, \Lambda \otimes A^{-1})\ \mathcal{IW}_P(\nu_0, V_0) \\ \times \mathcal{N}(\mathbf{log}\lambda_\mathbf{0}, \mathbf{\Sigma})\ \prod_t [\Phi(util_{it})]^{y_{it}}[1 - \Phi(util_{it})]^{(1-y_{it})}\end{aligned} \tag{1.11}$$

$$util_{it} = x_{it}\beta_{1_i} + d_{it}\beta_{2_i} + \tau_i\beta_\tau$$

The data suggest there is a good deal of asymmetry in reputation. Note that in Table 1.1, the standard deviation of cumulative uploads (1251.8) is more than double the mean (533.9). This indicates there is wide variation amongst publishers in uploading activity. Furthermore, the distribution is highly skewed, where 80% of the publishers have fewer than 604 cumulative uploads, yet the remaining 20% have on average 2010. This suggests that a sizable majority of publishers have significantly lower reputation than the top 20%. A positive coefficient on the interaction of expected share ($EShare_{it}$) and lawstock ($Law_{i,t-1}$) would indicate that when lawstock increases in the presence of competition, the publisher finds utility in increasing his activity to acquire reputation. However, a negative coefficient on the interaction would suggest the opposite. Our hypothesis is that publishers anticipate that increased lawstock may deter competition from acting to acquire reputation. This creates a "vacuum" which presents a larger share of voice opportunity for the focal publisher if he chooses to upload. We expect that low-reputed publishers' coefficients will be larger than their high-reputed counterparts. Since theoretically returns to reputation decline in current reputation level (Board and Meyer-ter Vehn, 2013; Resnick and Zeckhauser, 2002), gains of low-reputed publishers from acquiring reputation could be greater than that for high-reputed publishers. This would suggest that low-reputed publishers may have more incentive than high-reputed publishers to acquire reputation despite the increased costs due to increased law-stock. This concludes the description of the model. In the next section, we discuss endogeneity and identification of the model.

1.6.2 Endogeneity and Identification

One potential problem with endogeneity concerns the lack of data on downloads of the pirated content. We do not include expectation of content downloads in the publisher's utility because we do not observe the dynamics of downloading volumes on the platform. It is plausible that the publisher's decision to upload is partially driven by expected downloads of the content they publish. If downloads is an important omitted variable, it should have a correlation with not only the publisher's uploads, but also with the competitors' uploads. This would imply endogeneity concerns for the variables for reputation (Rep_{it}) and expected share ($EShare_{it}$), which are both functions of uploads. It is unclear whether consideration of size of audience is a key factor in the motivation to seek reputation-enhancement. Unlike traditional demand-supply models for priced goods where expected demand enters the utility function for supply, in our study reputation-building is based on the publishing of free digital goods. Digital goods can theoretically supply unlimited demand with no capacity constraints. Since the publisher's goal is to have his uploads seen and stand out from the crowd of competitors as effectively as possible, mere publishing of the content on the site exposes his name to platform participants. This might in itself have a positive impact on his reputation, regardless of how many people download his content.

Another potential endogeneity issue concerns whether copyright owner decisions to file lawsuits. It is plausible (in fact, likely) that the timing of lawsuit filings is not random and is instead driven by expectations of publisher uploading activity. For example, a software company may increase its filings in expectation of increased piracy of its latest release of its software application. Given that we model an individual publisher's decision to acquire reputation, it is not likely that copyright owners and trade organizations decide to file BitTorrent copyright lawsuits because of an individual publisher's decision to upload. This may somewhat relieve the severity of a potential endogeneity issue. Nevertheless, since much of the effects of lawsuits are common to all, our inclusion of daily fixed effects should presumably reduce endogeneity concerns about omitted variable bias during the estimation of the model (as well as for any other daily shocks common to all in the panel dataset).

To address potential endogeneity issues, we use the instrument-free copula method of Park and Gupta (2012). This method uses Gaussian copulas to model the joint distribution function between the error term and potentially endogenous regressors of Reputation (Rep_{it}), Lawstock (Law_{it}), and E[Share] ($EShare_{it}$). We include copula regressors in the publisher's utility function to account for the potential endogeneity issues outlined. The Gaussian copulas approach has the benefit that the task of finding valid instruments is unnecessary.

Because these potentially endogenous regressors have random coefficients in our probit model, the model structure rules out some potential alternatives for addressing endogeneity. Instrumental variable methods such as two-stage least squares or the control function method would be problematic to use in our model[4]. It is very difficult to find suitable instruments that are both strong and valid for our three candidate regressors. All three regressors inherently have a dynamic component to their nature, which conceptually results in past values being correlated to current values of the variable. Hence lagged versions of the variables are likely to also correlate with the error term and would therefore not make valid instruments. Furthermore, because the endogenous regressors have random coefficients, any instrument for the variable will be correlated with the error term, rendering instruments unsuitable for our setting (Ebbes et al., 2005; Wooldridge, 2015).

To avoid these issues, we turn to instrument-free methods that allow the researcher to model the correlation between the endogenous regressor and the error term. Two candidates we consider are Latent Instrumental Variables (Ebbes et al., 2005) and Gaussian copulas (Park and Gupta, 2012). Based on our tests and observations of the data, the regressors for Lawstock and Reputation appear to have an approximately negative-binomial distribution. The Eshare variable has a nearly exponential distribution. Latent Instrumental Variables (LIV) assumes that the en-

[4]Despite our concerns with IV methods, nevertheless we perform an analysis with IVs using (Lewbel, 1997) method of constructing instrumental variables. Discussion of these results can be found in Appendix D

dogenous regressor and the error term have a joint bivariate normal distribution. In our case, our data clearly violates this assumption. On the other hand, Park and Gupta's method relaxes that assumption and allows the endogenous regressors to take on any distributional form, as long as they are not normally distributed. The method assumes the error terms are normally distributed. If the endogenous regressors are also normally distributed, then identification issues arise. Given that our endogenous regressors are not normally distributed, identification should not present a problem. For these reasons, we address endogeneity by using Park and Gupta's method. This method is similar to control function methods in that we add as additional regressors the copula components that capture the correlation with the error term.

Model identification draws on the large unbalanced panel data of the daily publisher activity and from daily filings of BitTorrent copyright lawsuits over a four year period. Because we have daily activity of 752 publishers over a four year period, we have intra-temporal and cross-sectional variation in the focal publisher's uploads, upload sessions, and competitors' uploads. This rich variation provides the basis for identification of the parameters on the reputation, lawstock, and expected share variables.

Carry-over rates are identified from the speed at which uploading levels return to the pre-lawsuit amount of uploads a publisher typically publishes without the presence of lawsuits. Using similar logic, reputation carry-over rates depend on variation the focal publisher's uploading activity as well as their levels relative to competitive uploading activity. Reputation stock's associated carry-over rates are identified from the speed at which uploading returns to mean-level of uploads when there is zero-level of competition. The level 2 coefficients in the model (the prior on the random coefficients) are identified based on the variation due to time-invariant effects found in the panel data of the 752 publishers in the sample.

1.7 Estimation

We estimate the model using Bayesian Markov chain Monte Carlo (MCMC) methods. We first use Gibbs sampling (Gelfand and Smith, 1990) to estimate parameters of the rational expectations model (Equation 1.5) of each publisher's expected number of competitive uploads in the current period. Expected competitive uploads are estimated once and enter into the publisher's choice model as a fixed quantity. Conditional on the pre-estimated expected competitive uploads, we simultaneously estimate the mixed probit choice model, the reputation stock carryover rates, and lawstock carryover rates. We use a use a block Gibbs sampler with a Metropolis-Hastings (Chib and Greenberg, 1995) interim step to estimate the model in the following sequence:

1. Estimate the random coefficients portion of the level 1 mixed probit model, conditional on fixed effects and carryover rates.

2. Estimate the fixed coefficients portion of the level 1 mixed probit model, conditional on random coefficients and carryover rates.

3. Estimate level 2 model of the probit, conditional on random coefficients.

4. Estimate carryover rates for reputation and lawstock variables, conditional on all probit coefficients.

We use the Metropolis-Hastings (M-H) interim step in the last step of the sequence. Conditional on probit coefficients, the M-H algorithm compares log-likelihoods of the hierarchical probit model based on jointly drawn candidate carryover rates for reputation and lawstock. We parallelize the model on six CPU cores and run the model for 50,000 iterations, which are thinned to every fifth element to minimize auto-correlation. Trace plots, mean plots, and Gelman-Rubin (Gelman and Rubin, 1992) test results indicate that the chains converge (Gelman-Rubin = 1.03 - 1.05 for variables of interest). We burn 20% of the draws and use the remaining for the basis for estimation. We then sample from the joint posterior distribution to form an estimate of marginal statistics for the parameters of interest.

1.8 Analyses and Results

The estimated parameters that we present represent the average publisher's decision on whether to upload to build his reputation in the presence of competition. We compare our HB Probit model with a standard probit as well as a probit with publisher random intercepts. The parameters are summarized in Table 1.2. The estimates for the comparison models are in the first and second columns of Table 1.2, and our model for the paper is in the third column labeled "Paper's HB Probit'. The log likelihood, AIC, and BIC for the models suggest that our HB model has the best fit to the data. The publisher's current level of reputation (Rep_{it})

Table 1.2: Key Parameter Estimates of Publisher's Decision to Upload

	Dependent variable:		
	Pr(Upload)		
	Probit	*Random Intercept Probit*	*Paper's HB Probit*
Intercept	-3.949^{***}	-5.104^{***}	-6.473^{***}
	(0.021)	(0.041)	(0.471)
Reputation	-1.318^{***}	-1.603^{***}	-1.757^{***}
	(0.010)	(0.015)	(0.144)
Lawstock	-0.0003	-0.001	-0.130^{***}
	(0.002)	(0.003)	(0.033)
EShare	0.001	-0.071^{***}	-2.447^{***}
	(0.007)	(0.013)	(0.804)
EShare * Reputation	-0.020^{***}	-0.006^{***}	0.277^{***}
	(0.002)	(0.002)	(0.041)
EShare * Lawstock	0.004^{*}	-0.023^{***}	0.164^{***}
	(0.003)	(0.005)	(0.038)
Reputation * Lawstock	0.005^{***}	-0.002^{***}	0.057^{***}
	(0.0005)	(0.001)	(0.011)
Reputation Copula	6.274^{***}	8.299^{***}	9.055^{***}
	(0.036)	(0.052)	(0.385)
Lawstock Copula	-0.113^{***}	0.016^{*}	0.237^{***}
	(0.007)	(0.009)	(0.062)
EShare Copula	-0.070^{***}	0.034^{***}	11.447^{***}
	(0.004)	(0.006)	(3.077)
Daily FE	No	No	Yes
Observations	455,241	455,241	455,241
Log Likelihood	-98,911.5	-82,125.7	-76,386.5
AIC	197,843.0	164,273.4	155,440.9
BIC	197,953.3	164,394.8	155,682.9

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

has a negative effect on the decision to upload (-1.757, $p < 0.001$). The higher his current reputation level, the less utility he has to acquire more. This is consistent with findings from prior literature that utility to acquire more reputation decreases in reputation (Board and Meyer-ter Vehn, 2013; Resnick and Zeckhauser, 2002). The negative effect on reputation is also consistent with our proposed conceptual framework where we assert that publishers consider their current state of accumulated reputation when deciding whether to enhance it. Furthermore, reputation has a mean carryover rate (λ^R) and half-life of 0.494 and 0.98 days respectively ($p < 0.001$).

"Ceiling effects" pose a potential alternate explanation for the average decline in the utility for uploading as reputation increases. The presence of ceiling effects would imply that an exhaustion of available content inventory or restriction of available content could lead to declines in uploading activity. For an extreme example, a publisher specializing in TV may no longer have an incentive to upload content to the site if he has already uploaded every TV show ever produced in the world. To check whether there are any possible constraints on available content to publishers, we collected data on the number of titles available for sale on Amazon.com by conducting a search by content category on the site. Amazon sells entertainment products in all but one of the content categories on Kickass Torrents (the omitted category being XXX). Data from Amazon (specifically number of search results) provide a lower-bound baseline of the amount of entertainment content available, given that it is just one e-commerce site. Although it is the largest e-commerce site in the world, it is unlikely that it sells all the available movies, TV shows, games, etc. in the world. We collected Amazon search results from the website called the Internet Way Back Machine as of the final months of 2013 and compared that to the total number of uploads on Kickass Torrents in each category during the 2010 – 2013 period (see Table C.1 in Appendix C for details). The data suggest that total publisher uploads were only a fraction of the available content in each category on Amazon in 2013. Given that Amazon is only one of many potential sources of content, it is highly unlikely that publishers experienced ceiling effects on available content to upload.

Estimates on lawstock are consistent with our supposition that lawsuits are cost shocks to acquiring reputation on the piracy platform. The coefficient on lawstock (Law_{it}) is negative (-0.130, $p < 0.001$), which suggests that publisher i's awareness of lawsuits decreases the likelihood he would upload. Lawstock has a mean carryover rate (λ^L) of 0.567 ($p < 0.001$), implying a mean half-life of 1.22 days. Lawstock carryover rates have greater variation than reputation, where the standard deviation for the lawstock carryover rate is almost five times that of the reputation carryover rate (0.039 vs. 0.008). High variation in lawstock carryover rates could perhaps be due to high variation in the frequency that publishers are exposed to information about the lawsuits. Substantially lower variation in reputation carryover rates suggests that publishers on the platform may have greater collective consensus on the impact of one's reputation in the community. We did not find supporting evidence of significant correlation between the rates ($\rho = 0.06, p$ is n.s.).

The publisher's expectation of a share of voice gained from an upload ($EShare_{it}$)

is another factor in his utility for acquiring new reputation. Results indicate that a greater share of voice due to low levels of expected competition has a negative effect on the probability of uploading (-2.447, $p < 0.001$). This suggests that publisher i will be less motivated to upload unauthorized content when he expects small numbers of competitive uploads within a given time period than when there are hundreds of expected competitive uploads. Although the interaction of $EShare_{it}$ and Law_{it} is the focus of our study and provides the effect of interest, this result on the main effect of $EShare_{it}$ is unexpected. So we conducted deeper analysis to gain insight into the negative effect. We estimated the HB probit on 9 content categories of publishers. The are the content category most commonly associated with their uploads. Table A.1 in Appendix A displays the results of this analysis. We found that the effects of $EShare_{it}$ on the publisher's decision to upload is negative and significant for Movie publishers (-0.51, $p < 0.001$). It is negative, but statistically insignificant for Music and TV publishers (-0.01 and -0.15 respectively). For the remaining 6 content categories, the effect is positive, as expected, and statistically significant (Books: 0.62, $p < 0.001$; XXX: 0.52, $p < 0.001$; Applications: 1.06, $p < 0.001$; Anime: 0.40, $p < 0.001$; Games: 0.78, $p < 0.001$; Other: 1.30, $p < 0.001$).

Our current hypothesis that could explain these category-based results is that there could be a smaller selection of content to upload in Movies, Music, and TV compared to the other categories. A smaller selection of available content to upload could produce a "clumping" effect where publishers upload the same small set of available titles. Popularity of the titles could exacerbate this clumping effect. For example, imagine that only five movies were released in 2013 while 100 books were published in the same year. If there are only 5 movies available to pirate, then multiple publishers will upload the same five movies, which could produce a negative effect on the $EShare_{it}$ covariate during estimation. On the other hand, book publishers could me more widely distributed in the uploading of an available 100 book titles, which could produce a positive effect on the $EShare_{it}$ covariate. We conducted additional analysis on the Movie category to test this hypothesis by adding fixed effects for the top 5 and top 10 movies in terms of worldwide box office for each year of the data set. Although the size of the effect reduced in magnitude, the sign remained negative. Conducting additional analysis where all content titles were controlled for is analytically infeasible due to the thousands of titles relative to power in the dataset. Given that the main effect on $EShare_{it}$ is not the focus of this study, we leave additional inquiry on this interesting effect to future research.

Interactions between the key variables of reputation stock, lawstock, and competition allow us to answer the questions raised by this research: 1) do the effects of lawsuits (lawstock) impact publishers heterogeneously with respect their reputation levels? 2) does competition differentially impact publishers with respect to their reputation levels? 3) how does the publisher's decision change when he considers his competition with respect to lawstock levels? Our findings show that the coefficient on the interaction between expected share and lawstock is positive and significant (0.164, $p < 0.001$). This effect is consistent with our proposed conceptual framework: conditional on competition level, the more aware publisher i is of lawsuits, the more likely he is to act to acquire reputation. Lawstock may be an indicator of a possible reputation-enhancing opportunity, where lawsuits may drive out some competition and leave more opportunity to capture share of voice on the platform. Reputation's interaction with expected share and lawstock are also positive and significant (0.277, $p < 0.001$ and 0.057, $p < 0.001$ respectively).

We investigate whether there is any supporting evidence that heterogeneity in reputation levels of the publisher produces different outcomes in the decision to acquire reputation. We do this by examining whether the publishers with the lowest reputation levels in the community significantly differ from those with the highest levels of reputation in their utility for acquiring reputation. For our analysis, we segment reputation levels in terms of reputation stock per session (normalizing on a per session basis controls for tenure on the site). The top and bottom quartiles of these metrics define the low-reputed and high-reputed groups of publishers, respectively. Each of the groups contain 188 publishers (26% of the total). We use t-tests to compare the coefficients from the utility function to identify material differences between the groups. Table 1.3 summarizes the results. The general patterns we identify regarding how low-reputed publishers behave relative to high-reputed publishers are as follows:

1. Low-reputed publishers experience a smaller cost shock from lawstock (Law_{it}) than high-reputed (Low-Reputed: -0.092 vs. hi-reputed: -0.134, $p < 0.001$). This is consistent with our hypothesis that lawsuit awareness would be more costly to high-reputed publishers than low-reputed publishers.

2. Based on the interaction of $EShare_{it}$ and Rep_{it}, we find that low-reputed publishers have greater utility than high-reputed rivals to acquire reputation

Table 1.3: Low vs. High Reputation Publisher's Utility

	Low Rep	Hi Rep	Mean Diff.	
Intercept	-6.559	-4.298	-2.261	***
	(2.981)	(1.433)		
Reputation	-2.039	-0.533	-1.506	***
	(2.087)	(0.962)		
Lawstock	-0.092	-0.134	0.042	***
	(0.244)	(0.238)		
Eshare	-1.096	-1.166	0.070	***
	(1.103)	(1.103)		
EShare * Reputation	0.372	0.216	0.156	***
	(0.450)	(0.391)		
EShare * Lawstock	0.146	0.155	-0.009	***
	(0.305)	(0.297)		
Reputation * Lawstock	0.090	0.042	0.048	***
	(0.188)	(0.115)		
Reputation Copula	11.629	4.787	6.842	***
	(6.223)	(3.210)		
Lawstock Copula	0.232	0.214	0.018	***
	(0.435)	(0.427)		
EShare Copula	6.464	6.601	-0.137	***
	(3.544)	(3.545)		
Note:	*p<0.1; **p<0.05; ***p<0.01			

when they expect a greater share of voice from an action, conditional on their current reputations (Low-Reputed: 0.372 vs. hi-reputed: 0.216, $p < 0.001$). This is consistent with our conceptual framework that low-reputed publishers will be more motivated than high-reputed rivals to act to acquire reputation when competition is low.

3. The interaction of $EShare_{it}$ and Law_{it} indicates holding expected share of voice constant, low-reputed publishers are more likely than high-reputed rivals to act to acquire reputation when lawstock increases (Low-Reputed: 0.155 vs. hi-reputed: 0.146, $p < 0.001$). This is consistent with our hypothesis that low-reputed publishers are more likely to act to acquire reputation when expected share is favorable and lawsuit awareness increases in intensity.

4. The interaction of Rep_{it} and Law_{it} indicates that low-reputed publishers have greater utility to acquire reputation than high-reputed rivals as lawstock increases, conditional on their current reputation levels their utility function (Low-Reputed: 0.090 vs. hi-reputed: 0.042, $p < 0.001$)

This is consistent with our earlier proposal that low-reputed publishers are more motivated to upload to acquire reputation than high-reputed publishers when lawsuits are employed. This finding also suggests that lawsuit campaigns may be less effective in reducing piracy activity of low-reputed publishers than high-reputed publishers on the platform in the presence of competition for reputation.

1.9 Discussion

1.9.1 Summary

The goal of this research is twofold: 1) to study competitive reputation-building by consumers in a marketplace and 2) how competitive reputation-building can impact and be impacted by firm interventions (firm attempts to change consumer behavior). We use a novel dataset scraped from a popular online BitTorrent piracy platform called Kickass Torrents. This data captures uploading activity of 752 publishers over a 4 year period. We model the decision to upload unauthorized content in order to build their reputations as a recognized supplier of content on the platform. We also employ a second dataset of thousands of copyright lawsuits filed in the U.S. to investigate under what conditions the lawsuits change the dynamics of reputation-building and the balance of competition on the site. For decades, owners of copyrighted content (e.g., movie, TV, and music producers and artists, software developers) and trade associations (e.g., RIAA, MPAA, SDA) have used copyright lawsuits in attempts to reduce piracy of their content with the belief is that such campaigns will reduce piracy. Hence, we examine whether the firm intervention of lawsuits were ultimately effective in reducing piracy on the platform.

We employ a Hierarchical Bayesian probit model of the publisher's decision in each period of whether to upload unauthorized content. Marginal costs to acquire reputation, competition from other publishers for reputation, and his own reputation level affect the publisher's decision. During the period of our study, publishers experience a series of reputation cost shocks in the form of messaging about BitTorrent copyright lawsuits. We model the impact of reputation and lawsuits as stock variables and the publisher's expectations about competition levels using a rational expectations framework.

Our empirical results are consistent with the notion that the lawsuits seem to decrease the publisher's utility to upload unauthorized content, as intended. Yet, our results are consistent with the novel theory we propose. Our theory proposes that publishers find utility in actions that build reputation, but gains in reputation decrease in level of competition for it. Consistent with this, we evidence that publishers are more likely to upload unauthorized content when their anticipated competition for reputation is sufficiently low, illustrating the role that competition can play in the quest to build reputation. We propose that under certain conditions, low-reputed publishers are even more likely than high-reputed publishers to upload. Our theory proposes that because returns to reputation diminish in reputation level, lawsuits seem to be effective in inducing highly-reputed publishers to slow their uploading activity on the site, which creates a supply "vacuum" and conditions for reduced competition. As a result, those who are low-reputed and have higher utility to acquire reputation may perceive cost shocks like lawsuits as not threats, but rather as opportunities. Our empirical results are consistent with this notion.

1.9.2 Contribution

Our research provides theoretical and empirical contributions. First, we contribute a novel theory to the reputation literature. Our theory proposes that under certain conditions, competitive dynamics can have a significant role in reputation-building behavior in marketplaces. Individuals may compete with each other in the building of their own reputations, and their incentives are heterogeneous in their current reputation levels. The intensity of competition during the building of reputations can directly impact the returns to reputation. Furthermore, our theory also contributes to the incentives literature by demonstrating that competitive reputation building between agents (a competitive social dynamic) can change the impact of incentives. This change can lead to incentive backfire by altering the balance of competition which subsequently can produce behavior that is the opposite of what the incentives intended. The results we provide, derived from our novel datasets, lend empirical support to our theory and contribute to the reputation, incentives, and piracy literatures by providing empirical evidence of competition between individuals who are building their reputations.

1.9.3 Limitations and Opportunities for Future Research

We note that there are several limitations to our study, which offer avenues for future research. First, although we model the decision of whether or not to take a reputation-enhancing action within a given time period, our model does not account for how many actions should optimally be taken within that time period.

The reality is that often platform participants will choose to take more than one action within a time period to acquire reputation. Understanding the drivers behind optimal number of actions is a worthwhile pursuit for future research. Second, our research examines activities to acquire reputation that have no observable financial incentives. It is uncertain how financial incentives would change the dynamics, if at all, of competition for reputation, and whether the tensions between marketplace competition and competition for reputation move in the same or orthogonal directions. Third, we examine reputation acquisition in a black market where reputation is directly derived from illegal activity. Thus, the building of reputation-level, while positive in the eyes of the intended audience (downloaders) has negative repercussions for other stakeholders (copyright holders and law enforcement). As a result, the reputation-built simultaneously holds high value and high risk to the supplier. It would be interesting to examine whether the dynamics we find in our study hold under a context where possessing high reputation-level has positive value to intended audiences as well as to other stakeholders.

1.9.4 Conclusion

Our findings present some important implications for managers of platforms and online marketplaces with reputation systems. Using our findings as an example, the Kickass platform has a reputation system which transparently displays signals of reputation levels to all participants. Furthermore, the platform environment is one where suppliers on the platform compete for reputation, a highly sought-after asset.

Lawsuit campaigns, as a strategic mechanism, alter the balance of competition and under certain conditions can increase the targeted activity, especially among those of low-repute. Given that the majority of publishers on the platform have lower levels of reputation, lawsuit campaigns may have the effect of increasing piracy activity in the aggregate.

Platform managers and policy-makers seeking to influence participant activity on platforms or in marketplaces may want to consider design incentive mechanisms and policies that target reputation-enhancing activities. Furthermore, these mechanisms should take into account that 1) participants may compete with each other for reputation 2) participant utility and costs for acquiring reputation may be heterogeneous in reputation-level. For example, raising the costs to acquire reputation with the intention to reduce activity may not necessarily be an effective course of action. As our research demonstrates, in the presence of competition for reputation, increasing costs to a level may be sufficient in reducing the high-reputed segment of participants yet may not be sufficiently high enough to discourage the low-reputed segment. As a result, cost increases can actually encourage low-reputed platform participants to increase their activity. Thus, we recommend that managers carefully analyze the reputation seeking behaviors on their platforms and understand the costs and benefits that participants experience when they acquire reputation. By understanding this, managers can then set incentive mechanisms needed to manage agent behavior.

Chapter 2: Discrimination in Service

2.1 Introduction

In May 2017, the city of Philadelphia filed a lawsuit claiming that Wells Fargo & Company violated the U.S. Fair Housing Act. The lawsuit accused the bank of discriminatory lending practices that resulted in excessive foreclosures of minority-owned properties in the city (Phippen, 2017). Philadelphia's lawsuit came on the heels of similar suits by Oakland, Los Angeles, Miami, and Baltimore. Wells Fargo had already made a $175 million settlement with the U.S. Department of Justice in July 2012 for alleged discrimination against minority borrowers from 2004 to 2009 (Broadwater, 2012). Given its prior history, how was it possible that Wells Fargo was addressing service discrimination again in 2017?

Neither Wells Fargo nor the financing industry is unique on this issue. Nor does service discrimination concern only ethnicity or gender. Litigation concerning firm-against-consumer discrimination has a long history in the U.S. In the past two decades alone, prominent corporations such as Denny's, Walmart, Macy's, and Ally Financial paid more than half a billion dollars in settlements and fines for consumer discrimination cases based not only on race, ethnicity or gender (Elliott, 2003; Gutierrez, 2015; Koren, 2016), but also on age (Silver-Greenberg, 2012), geography

(Schroeder, 2017), social class (Kugelmass, 2016), and occupation (Addady, 2016). This amount does not include additional sales losses due to bad publicity, damaging boycotts, or impaired reputation and brand. Societal shifts motivate an increased need for more research on why discrimination in service still happens in the 21st century (Anderson and Ostrom, 2015; Bone et al., 2010; Hill and Stephens, 2003). Social fissures created by service discrimination have a direct impact on consumer and societal well-being (Bone et al., 2014; Crockett et al., 2011).

We define discrimination in service, or service discrimination, as different service treatment of consumers of equal quality (e.g., profitability to the firm) who differ only in group membership. We distinguish discrimination from prejudice in that prejudice, stereotypes, and racism focus on internally-held attitudes, beliefs, and ideologies. In contrast, discrimination, consistent with sociological and statistical discrimination literatures, concerns decision outcomes that exhibit unequal treatment of people based on the category to which they belong and is independent of internally-held attitudes (Arrow, 1973; Pager and Shepherd, 2008; Phelps, 1972; Quillian, 2006). Some service providers are undoubtedly prejudiced. However, suppose a firm has no prejudiced nor discriminatory intent, and its employees are not bigots? Under what conditions can service discrimination still emerge? What is the impact on profits over time? Extant literature does not provide a clear consensus on answers, indicating that the mechanisms producing service discrimination and its impact on profitability are still not well understood. Where prior literature primarily focuses on average attributes of consumer groups (e.g., average group wealth or average education) in explaining discrimination, we investigate how the variation in attributes (specifically, the distribution of consumer quality–defined as the true, latent profitability of the consumer) has a demonstrable effect on the magnitude of service discrimination. Furthermore, to the best of our knowledge, we are the first to investigate how social and competitive dynamics emerging from service discrimination impact long-term profits.

Our research employs a mixed-methods approach. We first present a theoretical model of how variance in consumer quality and magnitude of error in measuring consumer quality impact the magnitude of service discrimination. We validate the mechanism with empirical evidence from human decision makers. We then extend the analytical model with an agent-based model to investigate how competition and word-of-mouth emerging from service discrimination impact profits over time. We find that Group-Aware service decisions – decisions regarding which consumers to serve that are a function of the groups (i.e., ethnicity, gender, social class, education level) that consumers are members of – are perhaps profitable in the short-run. However, they can backfire over the long-run due to the effects of word-of-mouth and competition. Large measurement error in detecting consumer quality exacerbates service discrimination, while large variance in customer quality attenuates it.

The implications of this research are three-fold. First, it provides insight into how variance in consumer quality and measurement error can drive the emergence of discrimination, even if the firm is not inherently prejudiced. Second, it demonstrates how service discrimination's interaction with consumer word-of-mouth and competition damages profits. Third, this research suggests remedies that involve reducing error in measuring consumer quality or reducing the role that group membership information plays in assessing a consumer's quality (profitability). We elaborate on these themes in the remainder of the paper as follows: first we discuss how our research expands extant knowledge on discrimination in service. We then present our theoretical model and propose a definition of discrimination in service. We next present empirical validation of our theory. Then, we present an agent-based model (ABM) to go beyond the limitations of the analytical model: the ABM enables investigation into emergent macro-phenomena from the micro implications of the analytical model and studies. With the ABM, we investigate the long-term impact of service discrimination and its interaction with competition and word-of-mouth. Finally, we conclude with a discussion of the managerial and policy implications of our findings. In the next section, we discuss how this research contributes to the literature.

2.2 Extending the Literature on Service Discrimination

By extending the classical statistical discrimination model, we contribute to the literatures on the dynamic effects of discrimination in service. Our model provides insight on the macro effects on demand and profits which emerge from individual decisions that result in discrimination in service. We model both the firm-side dynamic decisions in its selection of consumers who apply for service as well as the consumer's decision of who to seek service from in response. With our research, we claim three contributions to the discrimination, marketing, and service literatures: First, to the best of our knowledge, we are the first to provide theory on the potential impact of service discrimination on consumer word-of-mouth and consequently on long-term demand for services. Second, while prior literatures such as Aigner and Cain (1977), Arrow (1973), Phelps (1972), and Lundberg and Startz (1983) give some treatment to the effects of variation within groups on the magnitude of discriminatory outcomes, their treatment of variance assumes that variance is static in nature. In contrast, we demonstrate the dramatic effects that dynamics in intra-group variation can have on the magnitude of discriminatory outcomes. Third, to the best of our knowledge, we are the first to integrate an analytical model with an agent-based model. The integrated models provide unique insights into the emergent macro effects over time of discrimination in service and constitute a methodological contribution. We further elaborate on these contributions in the following paragraphs.

Transformative consumer research (Mick, 2006; Pettigrew, 2001) and transformative service research (Anderson et al., 2010) contain extant literature on service discrimination from the consumer's point of view on topics such as financing options(Bone et al., 2014), consumer racial profiling (Crockett et al., 2003; Evett et al., 2013; Harris et al., 2005), and consumer discrimination against businesses selling "ethnic" French products in English-language dominated parts of Canada (Ouellet, 2007). In contrast, our work models service discrimination from the viewpoint of the firm. There are two primary theoretical camps in which research from the firm's viewpoint lies. The first, the taste for discrimination literature, theorizes that firms who discriminate may have a disutility for interacting with members of certain groups. They include this disutility in their objective function, which is not profit-maximizing (Becker, 1957; Schelling, 1969). A real-life example is a Colorado bakery which in 2012 refused to provide a wedding cake to a same-sex couple because of its religious-based service policy (Savage, 2017). In contrast, our research aligns with the second camp–the statistical discrimination literature (Aigner and Cain, 1977; Arrow, 1973; Coate and Loury, 1993; Phelps, 1972). This literature assumes firms are profit-maximizing and do not have a disutility for interacting with certain groups. Instead, the literature models discrimination as a problem of incomplete information where decision-makers use observable attributes such as group membership to draw inferences about individuals.

The bulk of prior models of service discrimination are static. We model discrimination dynamics, which is relatively sparse in the literature (Fang and Moro, 2011). Notable examples of dynamic discrimination research primarily examine the supply-side impact (labor and employment) of discrimination on the profit function (Antonovics, 2006; Bjerk, 2008; Blume, 2005, 2006; Bohren et al., 2017; Craig and Fryer, 2017; Fryer, 2007). Our study differs from these studies in the following important ways. They provide theory to explain implications of discrimination in employment, which are supply-side, cost-based human resource decisions. In contrast, we show the impact of discrimination in which consumers to serve on word-of-mouth, a demand-side, revenue-based sales decision.

Furthermore, extant literature primarily focuses on average group attributes (e.g., average wealth) and the assumption of fixed intra-group variation in the formation of the firm's beliefs about the individual and group. Few literatures have examined the impact of intra-group variation rigorously, and those that do assume that variance itself is constant. Aigner and Cain (1977), Arrow (1973), and Phelps (1972) proposed that groups with higher measurement error in detecting individual quality may experience greater discrimination due to uncertainty in the firms' assessment of their quality. Lundberg and Startz (1983) found that if two groups differ in their intra-group variation, there are conditions where social welfare can increase if firms ignore using group information in their quality assessment of employees. In contrast, we examine how changing intra-group variation over time (e.g., the distribution of group wealth) can impact firm beliefs and service discrimination. Because statistical discrimination is a problem of incomplete information, diagnosing the impact on discrimination that changing intra-group variation has, which introduces additional uncertainty, is critical to understand.

We also differentiate ourselves from prior literature in our methodological approach. There are only two papers we are aware of that use agent-based models along with other methodologies, however neither integrate an analytical model with an agent-based model. Peres and Van den Bulte (2014) use both an analytical model and agent-based model to examine when firms should use product exclusivity. However, the two models are not used in an integrated fashion. Instead, the analytical model is separately used to understand the generalizability of a parameter explored in the agent-based model. Smith and Rand (2017) integrate agent-based modeling with lab-based experimental studies to examine the macro-level, long-run implications of unemployed individuals using their social networks in job search. In contrast, we make a methodological contribution by integrating an analytical model with an agent-based model and experimental studies to examine the dynamics of service discrimination and its impact on demand and profits. As far as we know, we are the first to use an integrated analytical/ABM model, where the analytical model directly informs the ABM. In the next section, we present our theory of how variation in intra-group attributes has both immediate and long-term impact on the emergence of service discrimination.

2.3 A Model of Service Discrimination

Our theoretical model of service discrimination applies to a variety of service contexts that meet three criteria: 1) consumers can be segmented into two or more distinct groups based on an observable attribute; 2) service providers use both group membership based on the observable attribute and other information about the individual information to screen and determine level of service to the customer; 3) the service decision is conditional on a given price for service. Note that we examine the decision of whether or not to provide service to the prospective customer, which is a separate decision to be made before the pricing decision. We acknowledge that a service provider may choose instead to serve all prospective customers willing to pay a profit-maximizing price (such as a bank loan officer offering different interest rates to different loan customers). This is the equivalent of the firm first deciding to serve all customers with the willingness-to-pay, in which case price discrimination becomes relevant. That is not the scenario we assume here. For those interested in the topic of price discrimination, we refer readers to Bergemann et al. (2015), Narasimhan (1984), and Varian (1989) for excellent insights into this area.

Examples of services that *explicitly* screen potential customers include higher education, financial services, lodging, real estate, medical services, and membership-based services (e.g., country clubs). Examples of services where providers may *implicitly* screen customers include law enforcement, dining, hospitality, and retail. A recent example occurred in April, 2018 at a Philadelphia Starbucks where two African-American real estate brokers were arrested while waiting for a friend to arrive for a meeting. The store manager who called police claimed that the men were trespassing. In a video recording of the incident that went viral, other Starbucks customers can be heard protesting the arrests. White social media commentators wondered aloud why they had not been arrested for doing the exact same thing at Starbucks (waiting for friends before ordering). Subsequently, the incident generated substantial negative publicity, organized protests, a shutdown of the Philadelphia store, and a public apology media tour by the CEO of Starbucks who stated that "practices and training led to the bad outcome" (McLaughlin, 2018).

The initial setup of our model is consistent with the statistical discrimination theory of Nobel Prize winner Edmund Phelps (1972). We refer the reader to that paper and Aigner and Cain (1977) for details of the model and its derivations. However, the Phelps model does not offer insights that the impact of consumer heterogeneity, measurement error, and dynamics has on demand and ultimately on long-term impact on firm profits. We expand the model by exploring these areas. However, we first describe our model's initial setup, consistent with (Phelps, 1972), to provide a foundation from which we launch our contribution. Although our model applies to many service contexts, we ground our exposition in a specific bank lending context to facilitate intuition. We model a loan officer's decision regarding which applicants to give a bank loan. For simplicity, we assume that the bank has one loan officer and that the loan amount and interest rate are predetermined. The loan officer's decision is only whether to offer the loan to the applicant.

We also assume that the loan officer is boundedly rational (March and Simon, 1958): the loan officer uses locally available information (i.e., the historical credit scores of applicants at his bank only) and has finite cognitive and computational resources available to him in his decision-making. He updates his beliefs, based on past information, but he is not forward looking nor game-theoretic in his decision-making. He also does not have knowledge of his competitors' information or beliefs. We believe bounded rationality is a reasonable assumption because of what we learned from interviews with loan analysts at financial institutions. Their primary sources of information used in loan decisions include the applicant's credit score, credit history, current income and assets, and a fourth category they called "character factors". Character factors include any qualitative information loan analysts could find about the applicant's general character. For example, one interviewed loan analyst gave an example of a case where she found evidence of an applicant's history of gambling. Although the applicant had a sufficient credit score, income, and assets, he was not offered a loan. Interviewed loan analysts were consistent in stating that they used their institutional historical data from past applicants and loan recipients to compare current applicants in the decision-making process. They did not use information from competing financial institutions nor looked at future trends of applicant groups in making loan decisions for applicants.

Let us assume that each applicant i is a member of one and only one group $j \in \{H, L\}$, a high or low quality group. In our model, at the group level the mean quality, A_j, of the applicants within a group define whether the group is high or low. However, individual applicant quality can vary within a group. Quality is not a value judgement of the worthiness or status of the individuals in the group, but rather the average quality of the prospect of the consumers to the firm. We assume there is inequality between the groups, where

$$A_H > A_L > 0 \tag{2.1}$$

Examples of differences in quality between groups are upper versus lower social class, men versus women, racial, ethnic, or religious majorities versus minorities, college-educated versus less-educated, etc. For example, let us imagine that the applicants are segmented based on their residential address. The H applicants may live in an affluent part of town while the L applicants may live in a working-class area. We define inequality as $Inequality = A_H - A_L$.

Let us assume that the loan officer, conditional on a fixed interest rate[1], seeks to maximize profit by offering loans only to applicants whose quality, Q_{ij}, exceeds a threshold Q^{min}. A firm's decision of selecting consumers to provide service to, conditional on a given price, can be found in other contexts as well, such as lodging, insurance, financial services, medical services, etc. Quality (Q_{ij}) is a latent attribute of the consumer that is directly related to achieving the firm's objective function.

[1]We acknowledge that a bank officer might instead offer different interest rates to applicants of varying quality, in which case price discrimination becomes relevant. That is not the scenario we assume here.

In for-profit contexts such as this, quality can be interpreted as the applicant's true profitability to the firm. The threshold Q^{min}, assumed to be exogenous, is the quality of the marginal customer whose true profitability is 0 to the firm. We assume that Q_{ij} is normally distributed around group j's mean quality, A_j. Since the loan officer cannot observe Q_{ij}, the loan officer uses available information about each applicant (in the banking context, this can be credit history, net worth, income, debt, employment history, etc.) to form his expectation of the applicant's quality which subsequently informs his service decision. Information about the applicant is summarized in a single numerical score, S_{ij}, which is a noisy measure of Q_{ij}. The relationships between A_j, Q_{ij}, and S_{ij} are as follows:

$$Q_{ij} = A_j + \upsilon_{ij}, \quad \upsilon_{ij} \sim \mathcal{N}(0, \sigma^2_{q_j}) \tag{2.2a}$$

$$S_{ij} \mid Q_{ij} = Q_{ij} + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2_{\varepsilon_j}), \text{ where } \upsilon_{ij} \perp \varepsilon_{ij} \tag{2.2b}$$

$$S_{ij} \sim \mathcal{N}(A_j, \sigma^2_{q_j} + \sigma^2_{\varepsilon_j}) \tag{2.2c}$$

The loan officer wants to predict applicant quality, Q_{ij}, using the applicant's score, S_{ij}. Because S_{ij} has error, the loan officer supplements the score with information about the group of which applicant i is a member. He learns from the data he has on bank applicants and updates his beliefs about the mean quality of the groups with available information on applicant scores. Since A_j is latent, the loan officer forms a belief of the mean and variance of quality for each group, which he updates with the scores of other applicants who are members of the same group. Let us assume that the true mean and variance of quality of each group is unknown to the loan officer. But let us also assume that he knows that the mean is normally distributed and that the variance has an inverse-gamma distribution. These assumptions are consistent with Bayesian updating models of data that is normally distributed with unknown mean and variance (Gill, 2007). Let us assume that the loan officer has a jointly distributed normal-inverse-gamma prior on the mean and variance of quality of each group. These priors could be based on his experience with the bank, research conducted on the consumer market, or on data in an earlier period. Then his posteriors after Bayesian updating with score data are as follows:

(2.3)

$$P(A_j \mid \sigma^2_{q_j}, S_j) \sim \mathcal{N}\left(\frac{n_0 A_{j0} + n_j \bar{S}_j}{n_0 + n_j}, \frac{\sigma^2_{q_j}}{n_0 + n_j}\right)$$

$$P(\sigma^2_{q_j} \mid S_j) \sim \mathcal{IG}\left(\frac{n_0 + n_j}{2}, \frac{n_0 \sigma^2_{q_{j0}} + n_j \bar{\sigma}^2_{q_j} + \frac{n_0 n_j}{n_0 + n_j}(A_{j0} - \bar{S}_j)^2}{2}\right)$$

where

$$\bar{S}_j = \frac{1}{n_j}\sum_{i=1}^{n_j} S_{ij}$$

$$\bar{\sigma}^2_{q_j} = \frac{1}{n_j}\sum_{i=1}^{n_j}(S_{ij} - \bar{S}_j)^2$$

$$\{n_0, A_{j0}, \sigma^2_{q_{j0}}\} = \{\text{priors on number of data points, } A_j\text{, and } \sigma^2_{q_j}\text{, respectively}\}$$

$$n_j = \text{current number of data points}$$

Consequently, the mean of the posterior distributions in Equations 2.3 are the loan officer's expectations of the mean and variance of the quality of group j:

(2.4)

$$E(A_j \mid \sigma^2_{q_j}, S_j) = \hat{A}_j = \frac{n_0 A_{j0} + n_j \bar{S}_j}{n_0 + n_j}$$

$$E(\sigma^2_{q_j} \mid S_j) = \hat{\sigma}^2_{q_j} = \frac{n_0 \sigma^2_{q_{j0}} + n_j \bar{\sigma}^2_{q_j} + \frac{n_0 n_j}{n_0 + n_j}(A_{j0} - \bar{S}_j)^2}{n_0 + n_j}$$

We assume the loan officer again uses Bayes rule as well as the relationships established in Equation (2.2) to form his beliefs about the quality of applicant i from group j. His expectation of $Q_{ij} \mid S_{ij}$ is a weighted combination of information about the individual applicant (S_{ij}) and the loan officer's belief about the mean quality of the group she belongs to ($\hat{A}_j$). It can be conceptualized as a linear relationship with the following form:

$$Q_{ij} \mid S_{ij} = \gamma_j S_{ij} + (1 - \gamma_j)\hat{A}_j + \delta_{ij} \tag{2.5a}$$

$$E(Q_{ij} \mid S_{ij}) = \gamma_j S_{ij} + (1 - \gamma_j)\hat{A}_j \tag{2.5b}$$

$$\text{where } \delta_{ij} \sim \mathcal{N}\left(0, \gamma_j \sigma^2_{\varepsilon_j}\right) \text{ and } \gamma_j = \frac{\hat{\sigma}^2_{q_j}}{\hat{\sigma}^2_{q_j} + \sigma^2_{\varepsilon_j}} \tag{2.5c}$$

The quantity γ_j is known as the reliability of a measurement in classical score theory (Novick, 1965). It indicates how much reliance is placed on information about the individual applicant (score) relative to information about the group she's from to measure the targeted latent attribute. The score reliability has the following important properties:

$$0 < \gamma_j < 1, \qquad \frac{\partial \gamma_j}{\partial \hat{\sigma}^2_{q_j}} = \frac{\sigma^2_{\varepsilon_j}}{(\hat{\sigma}^2_{q_j} + \sigma^2_{\varepsilon_j})^2} > 0, \text{ and } \frac{\partial \gamma_j}{\partial \sigma^2_{\varepsilon_j}} = \frac{-\hat{\sigma}^2_{q_j}}{(\hat{\sigma}^2_{q_j} + \sigma^2_{\varepsilon_j})^2} < 0 \tag{2.6}$$

These properties highlight the impact that variation in quality ($\hat{\sigma}^2_{q_j}$) and score measurement error ($\sigma^2_{\varepsilon_j}$) have on the loan officer's beliefs about consumer quality. Increasing variation in consumer quality or decreasing variation in score measurement error increases the score reliability (γ_j). As score reliability increases, the loan officer places increasing weight on the consumer's individual information (S_{ij}) and less on group information ($\hat{A}_j$).

The top graph in Figure 2.1 (Tables and Figures follow the Reference section throughout) visually displays an example of the model using a range of 450 - 650 for S_{ij} (score) on the x-axis and a matching range on the y-axis for expected quality values, conditional on score: $E(Q_{ij} \mid S_{ij})$. The solid and dashed parallel lines are graphs of Equation (2.5b): the loan officer's expectation of quality of H and L applicants respectively. In this example, the two groups have the same score reliability, $\gamma = \gamma_H = \gamma_L = .5.$, where $\hat{A}_H = 723$, and $\hat{A}_L = 640$. A regression line coinciding

Figure 2.1: Two-Period Model of Service Discrimination

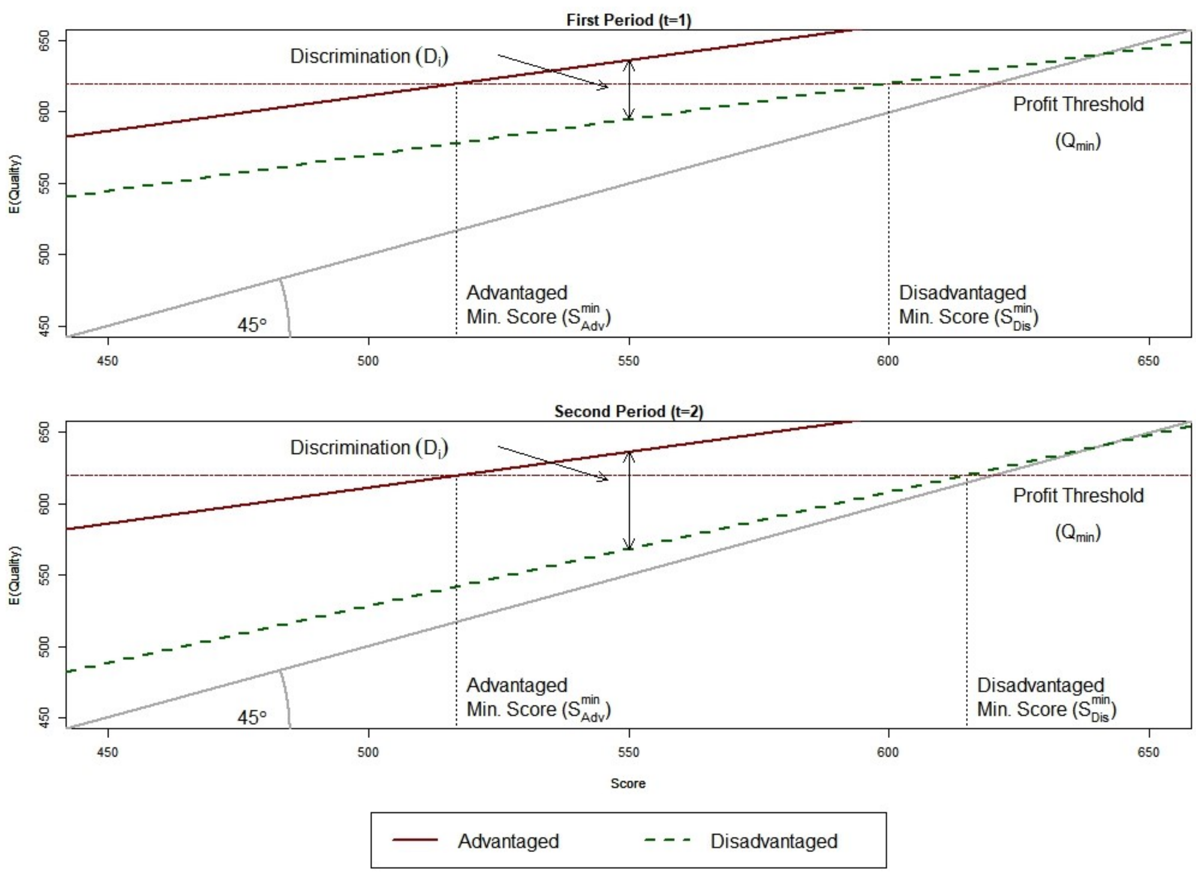

with the gray line at the 45° arc has a slope of $\gamma = 1$. This is where S_{ij} perfectly measures Q_{ij} without error. At this value, the loan officer has no need for group information $\hat{A}_j$ to form his expectation of customer i's quality. As measurement error is introduced, however, γ_j decreases. As $\gamma_j \to 0$, the regression representing the loan officer's expectation of applicant quality rotates clockwise towards a horizontal line with intercept $\hat{A}_j$. At its limit, $\gamma = 0$ and the customer's score S_{ij} no longer has weight. The loan officer has a monolithic belief about group j's members: $E(Q_{ij} \mid S_{ij}) = \hat{A}_j$.

The loan officer's expected profit from a single loan is

$$E(\pi_{ij} \mid S_{ij}) = E(Q_{ij} \mid S_{ij}) - Q^{min} \tag{2.7}$$

If the loan officer uses group information as well as the customer's score to form expectations about each applicant, then the loan officer uses a service policy where he offers a loan to applicants whose score, S_{ij}, meets or exceeds a minimum score criterion for their group. We subsequently refer to this service policy as the *Group-Aware* policy. We refer to the alternative policy of not using group information as the *Group-Blind* policy where all applicants face a single score criterion, regardless of group membership.

We derive the Group-Aware minimum score criterion for each group, S_j^{min}, by setting Equation (2.5b) equal to Q^{min} and rearranging terms.

$$S_j^{min} = Q^{min} + (Q^{min} - \hat{A}_j)\left(\frac{1-\gamma_j}{\gamma_j}\right) \tag{2.8}$$

In the top graph of Figure 2.1, the Group-Aware minimum score criteria of the example model are located at the vertical dotted lines labeled "**H** Min. Score (S_H^{min})" and "**L** Min. Score (S_L^{min})". Note that these vertical lines intersect with a horizontal dashed line labeled "Profit Threshold (Q^{min})" at the top right corner of the graph. Each intersection point is precisely where the expected quality of a member of the given group, conditional on score, is equal to the Q^{min} that represents the marginally profitable customer.

In contrast, the Group-Blind loan officer ignores group membership and aggregates all applicant information to form his expectations of quality and a single minimum score criterion, S_{all}^{min} (not shown on the graph). Because the errors associated with Q_{ij} and $S_{ij} \mid Q_{ij}$ are independent of each other (see Equation (2.2b)), aggregation of the two groups has no impact on σ_{ε}^2. However, aggregation does impact the mean and variation in customer quality. Let p_H ($p_L = 1 - p_H$) represent the proportion of all applicants that are members of the H (L) group. Using the equations for blended means and pooled variance, the mean quality, variance in quality, score reliability, and minimum score criterion of applicants under the Group-Blind policy is as follows:

$$\hat{A}_{all} = p_H \hat{A}_H + (1 - p_H)\hat{A}_L \tag{2.9a}$$

$$\hat{\sigma}_{q_{all}}^2 = \hat{\sigma}_q^2 + p_H(1 - p_H)(\hat{A}_H - \hat{A}_L)^2 > \hat{\sigma}_q^2 \tag{2.9b}$$

$$\gamma_{all} = \frac{\hat{\sigma}_{q_{all}}^2}{\hat{\sigma}_{q_{all}}^2 + \sigma_{\varepsilon}^2} > \gamma \tag{2.9c}$$

$$S_{all}^{min} = Q^{min} + \left(Q^{min} - \hat{A}_{all}\right)\left(\frac{1 - \gamma_{all}}{\gamma_{all}}\right) \tag{2.9d}$$

Let $f_j(S)$ and $F_j(S)$ represent the probability density function and cumulative distribution function of group j scores. The loan officer's expected profit (Π) under the Group-Aware and Group-Blind policies are respectively:

$$E(\Pi \mid S_{j\in\{L,H\}}^{min}) = \sum_{j\in\{L,H\}} \int_{S_j^{min}}^{\infty} \frac{p_j \; E(Q_{ij} \mid S_{ij}) f_j(S) dS}{2 - F_H(S_H^{min}) - F_L(S_L^{min})} \tag{2.10a}$$

$$E(\Pi \mid S_{all}^{min}) = \sum_{j\in\{L,H\}} \int_{S_{all}^{min}}^{\infty} \frac{p_j \; E(Q_i \mid S_i) f_j(S) dS}{2 - F_H(S_{all}^{min}) - F_L(S_{all}^{min})} \tag{2.10b}$$

Under conditions of incomplete information about true consumer quality, thus far the loan officer has taken a profit-maximizing, non-prejudiced approach to forming a service policy. So where is the discrimination in service? We formalize our definition of service discrimination (D_i) as follows:

Definition 2.1 *Discrimination in service occurs when the service provider differentially treats two equally qualified consumers (i.e., consumers with the same quality and score) just because they are members of different groups. It is equivalently defined as the service provider's change in treatment of consumer i if consumer i changes group membership, conditional on maintaining the same quality and score. Discrimination (D_i) is defined as*

$$\begin{aligned} D_i &= E(Q_{i,H} \mid S_i^*, Q_i^*) - E(Q_{i,L} \mid S_i^*, Q_i^*) \\ &= (\gamma_H - \gamma_L)Q_i^* + \left[(1-\gamma_H)\hat{A}_H - (1-\gamma_L)\hat{A}_L\right] \\ &\text{where } S_i^* = S_{i,H} = S_{i,L} \text{ and } Q_i^* = Q_{i,H} = Q_{i,L} \end{aligned} \tag{2.11}$$

The top graph in Figure 2.1 shows by example the magnitude of discrimination for consumers with a score $S_i^* = 550$. From Equation (2.11), we can see that if $\gamma_H \neq \gamma_L$, there are consumers with quality level $Q_i^* = Q_{D0} = \frac{(1-\gamma_L)\hat{A}_L - (1-\gamma_H)\hat{A}_H}{(\gamma_H - \gamma_L)}$ who experience no service discrimination. However, other consumers with quality levels higher or lower than Q_{D0} experience discrimination at magnitudes that increase in absolute value the further quality is from Q_{D0}. However, if $\gamma = \gamma_H = \gamma_L$, then the magnitude of discrimination is constant across all consumers. Discrimination in Equation (2.11) then simplifies to $D_i = (1-\gamma)\left(\hat{A}_H - \hat{A}_L\right)$.

While Phelps (1972) discusses the implication that his model implies that each group will have its own service threshold to receive service, we confirm and extend the Phelps model in the following findings that provide additional insights (see E for proofs):

Result 2.1 *A profit-maximizing service provider using a Group-Aware policy will set the same minimum score criterion for each group only if* $\gamma_H = \gamma_L \left[\frac{(Q^{min} - \hat{A}_H)}{Q^{min} - [\gamma_L \hat{A}_H + (1-\gamma_L)\hat{A}_L]} \right] = \gamma^*$*. Otherwise, the service provider will set a higher minimum score criterion for the L group than the H group* ($S_L > S_H$) *for all* $\gamma_L \in (0, 1)$ *when:*

1. $\gamma_H \in (0, \gamma^*)$ *and* $Q^{min} < \left[\gamma_L \hat{A}_H + (1 - \gamma_L)\hat{A}_L \right]$

2. $\gamma_H \in (\gamma^*, 1)$ *and* $Q^{min} > \left[\gamma_L \hat{A}_H + (1 - \gamma_L)\hat{A}_L \right]$

Under these conditions, L customers must meet a higher service policy criterion than H customers to receive the same level of service.

Consistent with prior statistical discrimination models, in general the service provider has a higher service criterion for L group consumers than for H group customers. What is interesting about this result, however, is that there are regions of the state space where the reverse is true: that it is profit-maximizing for the service provider to set $S_L < S_H$. The three charts in Figure 2.2 provide examples of the relationship between S_L and S_H in the state space. On the x-axis is a range of possible Q^{min} values. The y-axis represents values for γ_H. Each chart is distinguished by low (0.24), medium (0.50), and high (0.76) values of γ_L for top, mid, and bottom charts respectively. The dark blue regions in the charts represent conditions where the service provider would set a higher criterion for the L group than the H group. The white regions represent conditions where the service provider would set the reverse: a higher minimum score criterion for the H group. The black curves on the

Figure 2.2: Minimum Service Criteria State Space

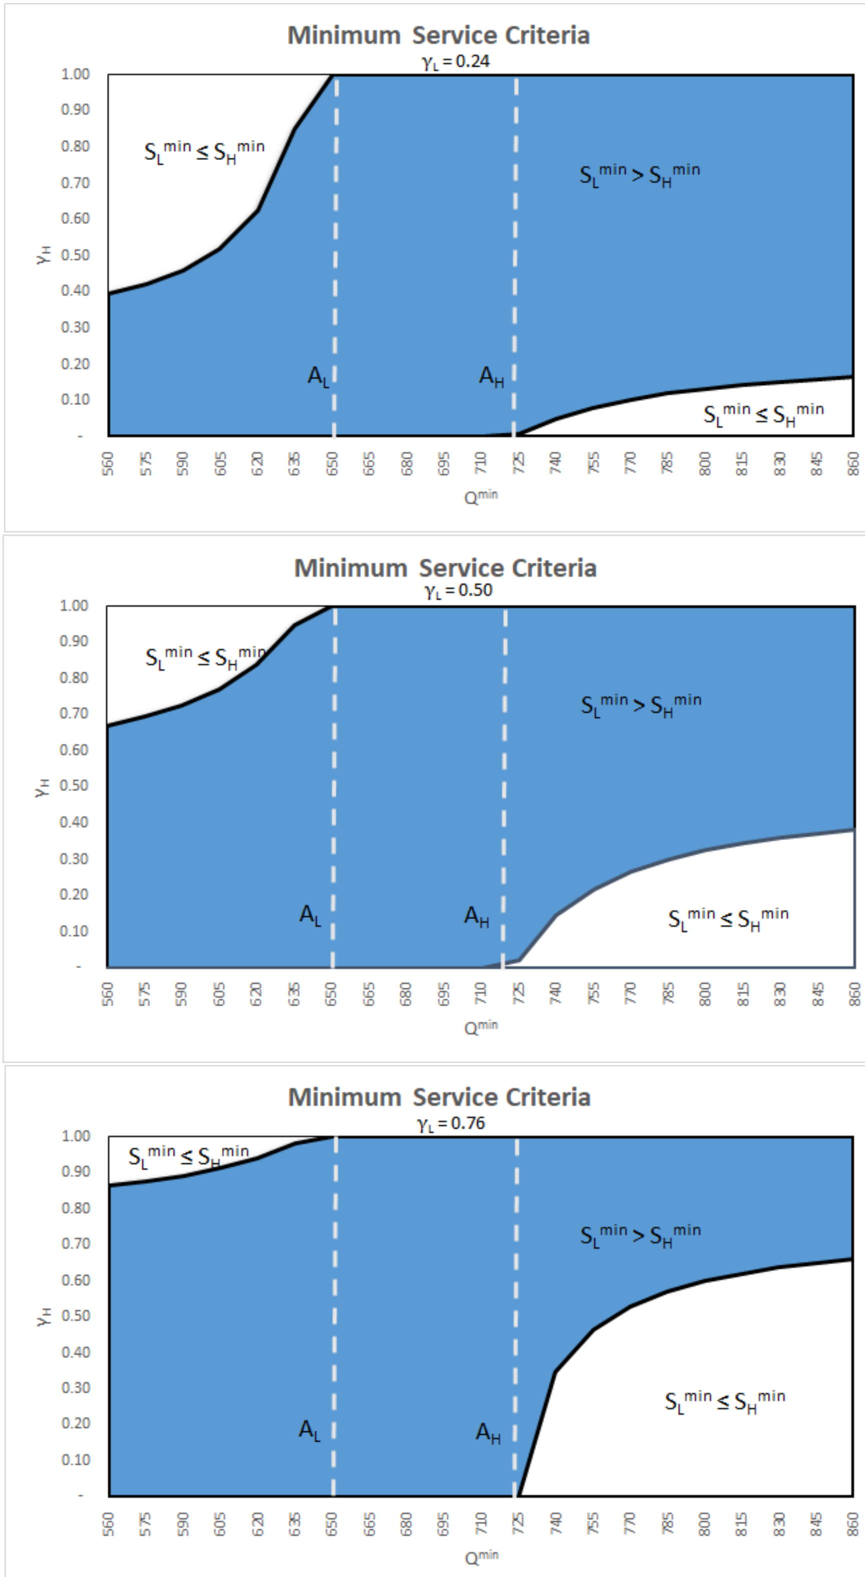

boundaries between the dark and white regions represent (narrow) regions where both groups have equal minimum service criteria. The light dashed vertical lines mark the values of $\hat{A}_L = 640$ and $\hat{A}_H = 723$ in these examples. Note that $S_H < S_L$ for all γ_j when $\hat{A}_L < Q^{min} < \hat{A}_H$. The condition indicates that a loan to an average L consumer is not profitable to the bank, but a loan to an average H customer is. However, $S_L < S_H$ only if $Q^{min} < \hat{A}_L$ or $Q^{min} > \hat{A}_H$. In these regions, consumers from both groups are on average profitable to the bank ($Q^{min} < \hat{A}_L$) or are both unprofitable to the bank ($Q^{min} > \hat{A}_H$).

Consequently, we find that it can be profitable to discriminate. The average per period (short-term) profits are greater from a Group-Aware service policy than from a Group-Blind one:

Result 2.2 *Let there be two consumer groups represented by* $j \in \{H, L\}$ *where one group's mean quality is higher than the other* ($\hat{A}_H > \hat{A}_L$)*. Also, assume that the service provider aggregates all customer information to set a single, group-blind minimum score criterion policy* S_{all}^{min}*. Furthermore, assume that aggregation has no impact on variation in measurement error,* σ_ε^2*. Then* $E(\Pi \mid S_j^{min}) \geq E(\Pi \mid S_{all}^{min})$*. In other words, the service provider's expected average per period profits from a Group-Aware service policy is greater than one using a Group-Blind service policy.*

It is important to note here that this finding is conditional on the assumption that the focal firm is a monopoly. As prior literature has suggested (Becker, 1957), allowance for competition can reverse this effect. We reaffirm this assertion under an alternative competitive context involving consumer word-of-mouth in our agent-based model, which we further elaborate upon later in this paper.

Taking the derivative of D_i in Equation (2.11) with respect to σ_ε^2, $\hat{\sigma}_q^2$, and γ yields:

Proposition 2.1 *Assume each group $j \in \{H, L\}$ has equal σ_ε^2, σ_q^2, and γ. Discrimination (D_i) varies in σ_ε^2, σ_q^2, and γ as follows:*

1. *D_i increases in σ_ε^2, the magnitude of error in measuring true customer quality.*

2. *D_i decreases in σ_q^2, the variation of customer quality within each group.*

3. *D_i decreases in γ, the reliability of individual information (score) about each customer.*

The intuition behind this proposition is that the greater the inequality between two groups, the greater the difference will be in the loan officer's assessment of two equally qualified customers from each group. However, the greater the reliability of individual customer information is in assessing quality, the less the loan officer will rely on group information to form his expectation. Reliability of individual information improves when there is more information about members within a group (i.e, more intra-group variation in customer quality) and when there is decreased error in measuring the quality of individuals. Increased reliance on the customer's score/decreased reliance on group information leads to decreased discrimination.

The discrimination model described thus far is static, similar to most prior literature. However, consumer attributes can change over time. In the U.S., for example, a woman's median hourly earnings was 64% of a man's in 1980. As of

2015, it was 83%. This pattern is consistent for many traditionally L groups in the US (Brown and Patten, 2017). For this reason, it is important to understand how dynamics in group inequality can affect the dynamics of service discrimination. Next, we consider a time dimension in our model, which is a departure from standard statistical discrimination models such as Phelps (1972) and Aigner and Cain (1977).

To explore dynamics, we expand the time frame to two periods and add the subscript t to the model (i.e., we now have A_{jt}, Q_{ijt}, S_{ijt}, $\sigma^2_{q_{jt}}$, $\sigma^2_{\varepsilon_{jt}}$, $\gamma_{jt}, S^{min}_{jt}$, and D_{it}). Let us assume that each group has a new set of applicants at $t = 2$ with a mean quality level that may differ from the mean quality level of applicants from $t = 1$. Let $p_{j1} + p_{j2} = 1$, where p_{j1} is the proportion of all group j applicants across time periods that applied at $t = 1$ and with mean quality level A_{j1}. The proportion that applied at $t = 2$ is $p_{j2} = 1 - p_{j1}$ with mean quality level $A_{j2} = g_j A_{j1}$. The value $g_j \in [0, \infty)$ is the magnitude of change in group j mean quality from $t = 1$ to $t = 2$. A $g_j = 1$ implies no growth in group j's mean quality between periods 1 and 2, $g_j < 1$ implies a decline, and $g_j > 1$ suggests growth. In the example about women's pay relative to men's over the last three decades, $g_{women} > 1$. For simplicity's sake, also assume that each group's new set of applicants has the same magnitude of intra-group variation in quality and variation in measurement error as their predecessors at $t = 1$ (i.e., $\sigma^2_{q_{j1}} = \sigma^2_{q_{j2}}$ and $\sigma^2_{\varepsilon_{j1}} = \sigma^2_{\varepsilon_{j2}}$).

If the loan officer uses score information from just the new set of applicants in $t = 2$ to form expectations of quality, the only change in his belief relative to $t = 1$ is driven by g_j, the change in group mean quality. There is no comparative change in variation in quality and measurement error between $t = 1$ and $t = 2$, therefore

$\gamma_{j1} = \gamma_{j2}$. On the other hand, if the loan officer Bayesian updates his expectations, then he may use what he learned about group j in the first period $(\hat{A}_{j1}, \hat{\sigma}^2_{q_{j1}})$, as his prior for the formation of expectations for the second period. Using Equations 2.3, the loan officer's posterior beliefs of the distribution of group j's current quality as of $t = 2$ is as follows:

Lemma 2.1 *Assume there are time periods $t \in \{1, 2\}$. If a service provider updates his expectations of group j's quality in a Bayesian fashion, he may use past information about group j as his prior on the beliefs of the current period. The following is true about his current beliefs of j's mean quality level, variation in quality, score reliability, and the minimum score criterion (where subscript c represent current beliefs as of $t = 2$):*

$$\begin{aligned} \hat{A}_{jc}(g_j) &= p_{j1}\hat{A}_{j1} + p_{j2}\hat{A}_{j2} \\ &= \hat{A}_{j1}\left[p_{j1} + g_j(1 - p_{j1})\right] \end{aligned} \tag{2.12a}$$

$$\begin{aligned} \hat{\sigma}^2_{q_{jc}}(g_j) &= p_{j1}\hat{\sigma}^2_{q_{j1}} + p_{j2}\hat{\sigma}^2_{q_{j2}} + p_{j1}p_{j2}\left[\hat{A}_{j1} - \hat{A}_{j2}\right]^2 \\ &= \hat{\sigma}^2_{q_{j1}} + p_{j1}(1 - p_{j1})\left[\hat{A}_{j1}(1 - g_j)\right]^2 \geq \hat{\sigma}^2_{q_{j1}} \end{aligned} \tag{2.12b}$$

$$\gamma_{jc}(g_j) = \frac{\hat{\sigma}^2_{q_{jc}}(g_j)}{\hat{\sigma}^2_{q_{jc}}(g_j) + \sigma^2_{\varepsilon}} \geq \gamma_{j1} \tag{2.12c}$$

$$S^{min}_{jc} = Q^{min} + \left(Q^{min} - \hat{A}_{jc}(g_j)\right)\left(\frac{1 - \gamma_{jc}(g_j)}{\gamma_{jc}(g_j)}\right) \tag{2.12d}$$

From Equations (2.12b) and (2.12c), we see that as long as there is a change in a group's mean quality from $t = 1$ to 2 (i.e., $g_j \neq 1$), then j's intra-group variance in quality and j's score reliability increases. These increases can either decrease or increase the group's minimum score criterion. If a group's mean quality increases, it is intuitive that its minimum score criterion would decrease. After all, increasing group quality implies that more members are qualified (profitable) to the service provider for services. It also intuitively follows that decreasing group quality should lead to an increasing minimum score criterion. However, we also find conditions where a group's minimum score criterion can increase (decrease) despite its increasing (decreasing) mean group quality. This is formalized in the next proposition:

Proposition 2.2 *The effect of growth or decline (g_j) of a group's mean quality over time on the group's most recent minimum score criterion, S_{jc}^{min} depends on the group mean quality levels relative to Q^{min}. When a group's mean quality changes over time (i.e.,$g_j \neq 1$), its minimum score criterion changes in the following ways:*

(2.13)

$$S_{jc}^{min} \begin{cases} > S_{j1}^{min} \text{ when } \hat{\sigma}_{q_{jc}}^2(g_j) < (\hat{\sigma}_{q_{jc}}^2)^* \text{ and } \begin{cases} g_j & > 1 \text{ and } Q^{min} < \hat{A}_{jc}(g_j) \\ g_j & < 1 \text{ and } Q^{min} > \hat{A}_{jc}(g_j) \end{cases} \\ = S_{j1}^{min} \text{ when } g_j = 1 \text{ or } Q^{min} = \hat{A}_{jc}(g_j) \\ < S_{j1}^{min} \text{ otherwise} \end{cases}$$

$$\text{where } (\hat{\sigma}_{q_{jc}}^2)^* = 2p_{j1}(1-p_{j1})\hat{A}_{j1}(1-g_j)(Q^{min} - \hat{A}_{jc}(g_j))$$

This means that when intra-group variation is not too high, a group can face unchanging or even rising minimum score criteria for service despite the group's improving quality. The intuition behind proposition 2.2 is that as score reliability improves, there is less uncertainty regarding the score's measurement of true quality. Since scores are more reliable, consequently the service provider choses a minimum score criterion that draws closer to Q^{min}. At its limit where $\gamma_j = 1$ which implies perfect score reliability, $S_j^{min} = Q^{min}$. If $\hat{A}_{jc} > Q^{min}$, any improvement in score reliability reduces the amount of adjustment needed, which thereby increases the minimum score criterion towards Q^{min}. On the other hand, if $\hat{A}_{jc} < Q^{min}$, any improvement in score reliability decreases the adjustment needed and thereby reduces the minimum score criterion.

If the two groups differ in their rates of growth in mean quality, then there are additional implications on discrimination. As is the case of men and women's growth in pay rates over the past three decades, let us assume that $g_H < g_L$: the L group has a faster growth rate in mean quality than the H group. We established from Equation (2.11) that if a group increases its mean quality, the variance of quality of its members increases. Increased variance in the L group implies that $\gamma_{L,2} > \gamma_{H,2}$ (recall that the groups had equal score reliability at $t = 1$). This has the effect that the loan officer places greater weight on individual information for the L group than he does for the H. The regression for the L group rotates closer towards the 45° line where S_{ij} perfectly measures Q_{ij}.

The bottom graph in Figure 2.1 displays an example of the effect of growth in mean quality for the L group. The bottom graph represents the service provider's expectations at $t = 2$. The H group has remained unchanged relative to $t = 1$ (the top graph). However, the L group has grown in quality, resulting in a change in score reliability from a $\gamma_{L,1} = .5$ in $t = 1$ to a $\gamma_{L,c} = .8$ in $t = 2$. The minimum score criterion for the L group has subsequently increased by 15 units to 615 in $t = 2$.

This change in the score reliabilities of each group due to changes in variation has a demonstrable effect on the magnitude of discrimination against the L group.

Although the L group has grown in quality, the inequality gap shrinks for only some of its members. Where $Q_{\Delta D=0}$ is the quality level where the magnitude of $D_{i1} = D_{i2}$, L group applicants with quality $Q_{i,L}$ on one side of $Q_{\Delta D=0}$ actually experience greater magnitudes of discrimination than they did in the prior period. L group members on the other side of $Q_{\Delta D=0}$ benefit from their group's growing quality–discrimination is less than its magnitude at $t = 1$. This insight brings us to our final proposition.

Proposition 2.3 *Let* $\Delta\gamma_j = \gamma_{j2} - \gamma_{j1}$ *and* $\Delta\hat{A}_j(1-\gamma_j) = \hat{A}_{j2}(1-\gamma_{j2}) - \hat{A}_{j1}(1-\gamma_{j1})$. *A change in the mean quality of a group can change the degree of discrimination that members of the L group experience over time. The quality level for which the magnitude of discrimination is the same in both time periods,* $Q_{\Delta D=0} = \frac{\Delta\hat{A}_L(1-\gamma_L)-\Delta\hat{A}_H(1-\gamma_H)}{\Delta\gamma_H-\Delta\gamma_L}$, *the magnitude of discrimination that customer i experiences at $t = 2$ can be less or greater than at $t = 1$. The conditions where the magnitude of discrimination changes over time t has the following relationship:*

1. $D_{i2} > D_{i1}$ *if*

   (a) $\Delta\gamma_H > \Delta\gamma_L$ *and* $Q_{ij} > Q_{\Delta D=0}$

   (b) $\Delta\gamma_H < \Delta\gamma_L$ *and* $Q_{ij} < Q_{\Delta D=0}$

2. $D_{i2} < D_{i1}$ *if*

   (a) $\Delta\gamma_H > \Delta\gamma_L$ *and* $Q_{ij} < Q_{\Delta D=0}$

   (b) $\Delta\gamma_H < \Delta\gamma_L$ *and* $Q_{ij} > Q_{\Delta D=0}$

Proposition 2.3 demonstrates that reducing the inequality gap does not reduce discrimination for everyone. When the L group grows in mean quality faster than the H group, its members with sufficiently high quality level can experience reduced discrimination over time. However, L group customers of sufficiently low quality can actually experience greater discrimination over time, even though the inequality gap is shrinking. Proposition 2.3 also asserts the reverse. Paradoxically, a growing inequality gap can reduce the magnitude of experienced discrimination for some L group members. When the inequality gap grows (the growth in quality for the H group exceeds that of the L group), L group customers of sufficiently low quality can experience reduced discrimination. The bottom graph in Figure 2.1 displays an example of the effect of Proposition 2.3. Note that the magnitude of discrimination for consumers with a score $S_i^* = 550$ is demonstrably greater than it is for them in $t = 1$ (top graph).

This completes our analytical model of service discrimination, which provides insight on how variation in customer quality and measurement error and have direct impact on the magnitude of discrimination in service. It also highlights how changes in group quality over time can change not only the variation in the quality of the group, but also the magnitude of discrimination. However, we now need to understand whether our analytical model is empirically realistic and valid. In the next section, we test the results from our analytical model with a set of empirical studies to understand if real behavior is consistent with our theory.

2.4 Empirical Validation

To test the empirical validity of our theory of service discrimination, we designed two studies analytical results presented in the prior section. For our first study, we recruited 400 participants (56.4% male, 5.5% between 25-34 years old, all U.S. residents) on Amazon mTurk. Study 1 tests the first three findings under conditions where the H group and L groups are constant in mean quality over time. For our second study, we recruited 409 participants on Amazon mTurk (54.0% male, 45.7% between 25-34 years old, all U.S. residents). In Study 2, mean quality is held constant for the H group while we allow the L group to grow monotonically in mean quality (approximately 1.36% after each round) during the study. By the last round, both groups are equal in mean quality. This condition enables us to test the last two propositions regarding dynamic inequality. This specification implies that $\gamma_L > \gamma_H$ over time, which meets the conditions of Proposition 2.3.1 (b) and Proposition 2.3.2 (b).

Each study had a 2 (group information: present vs. absent) X 2 (true quality variation: low vs. high) X 2 (credit score error variation: low vs. high) between-subjects design. Participants were randomly assigned to one of the eight conditions. We conducted the studies using Qualtrics survey software customized with JavaScript programming to dynamically update stimuli based on participants' prior decisions. Each participant played the role of a bank loan officer in 10 rounds of a lending game, which took on average 40 - 60 minutes to complete. The participant reviewed 10 loan applicant profiles for each of 10 rounds. Each profile revealed the applicant's identification number, address, and credit score. All applicants are members of one of two groups: either a square group or triangle group. We used these abstract symbols to represent groups in order to remove and explicitly rule out taste-based discrimination as the driver of the effects. Unknown to participants, one applicant group was the H group whose average quality and credit scores were higher than the other group. We counterbalanced between subjects the designation of whether the squares or triangles were the H group in order rule out any participant preferences for shapes or any taste-based discriminatory choice patterns in the study. In the Group-Aware (Group-Blind) condition, the study participant saw (never saw) of which group the applicant was a member.

The game asked the participant three questions about each applicant profile: 1) whether or not to offer the applicant a loan; 2) how confident the participant is in her decision; and 3) how confident the participant is that the loan will be repaid by the applicant. At the end of each game round, the game presented the participant a table displaying outcomes of the participant's loan decisions. The participant earned 30 Credit Units ($.075) on any repaid loans and lost 30 on loan defaults. After conclusion of the 10-round game, the participant answered survey questions about demographics and risk-aversion. The participants also earned a flat $3 fee for completion of the game. In total, we spent almost $10,000 on the two studies (including compensation to participants, pretests, software development, and Amazon fee). We gave relatively large mTurk compensation to ensure that participants took the task seriously and to test true decision-making behavior with financial consequences. Furthermore, we wanted to ensure participants were properly motivated and compensated to complete this lengthy study.

Randomly generated credit scores and loan repayment outcomes in the study were based on 2006 and 2010 Equifax data on mortgage applicants (Bhutta and Canner, 2013). The credit scores have a range of 300 - 850 in possible values, consistent with retail credit markets. We set $Q^{min} = 620$ (unobserved by participants), which is the value employed in the Equifax Risk report. Note that Q^{min} is below the mean quality levels of both groups in the studies ($\hat{A}_H = 723$ and $\hat{A}_L = 640$). Between study conditions, we manipulated the variance in applicant quality ($\hat{\sigma}^2_{qt}$) and measurement score error ($\sigma^2_{\varepsilon t}$) to be high or low values[2]. Participants did not see these values, but they directly influenced the credit scores on each applicant profile. Within each study condition, squares and triangles had the same intra-group variance in quality and measurement error. All participants in a study condition saw the same sequence of profile information. Our goal was to ascertain the impact of these conditions on our three dependent variables: probability of offering a loan to an applicant, confidence that the applicant will repay the loan (measured on a 0 - 100 point scale with 100 implying complete confidence), and profit made by the participants on loans offered. Next, we discuss how we test and measure the data collected to determine empirical validity of our theory.

2.4.1 Tests and Measures

Because the empirical study has repeated measures of participant response, we analyzed the data using hierarchical models to address likely intra-participant

[2]standard deviation=45.5 (Low) vs. 80.2 (High), based on the Equifax data

correlation of responses. Level 1 unit variables were study conditions and attributes of applicant profiles while participant dummies were at level 2. We used an OLS regression to test Result 2.2 because the dependent variable was simply total profits earned by each study participant. Dependent variables were a dummy indicator for whether the participant offered a loan to an applicant (*Loan*) and his confidence that the applicant would repay a loan (*Repay*, based on a 100-point scale). Study condition variables were dummy indicators for the Group-Aware condition (*Group*), low variation in applicant quality condition ($QualityLowVar$), low measurement error condition ($ScoreLowVar$), and Study 2: Dynamic Inequality (*Dynamic*). We also included continuous variables for study round (*Round* and $Round^2$). Applicant profile variables were applicant quality, score, and group membership (*Quality* $Score$, *LGroup* dummy).

To test Result 2.1 and Proposition 2.2, we modeled participant p's decision to offer a loan to applicant i from group j with a hierarchical logit.

$$\begin{aligned} Logit[Loan]_{pij} &= X_{pij}\beta_{10} + \beta_{0p} + \epsilon_{pij} \\ \beta_{0p} &= \upsilon_{00} + \xi_{0p} \end{aligned} \tag{2.14}$$

Although we did not observe each participant's implicitly determined minimum score criterion for applicants, we reason that if there is a higher score criterion in operation for L group applicants, we should observe a lower likelihood of loans offered to L group applicants than to equally qualified H group applicants.

To test Propositions 2.1 and 2.3, we estimated the following hierarchical linear model. The dependent variable is *Repay*, which is participant p's response to the question “How confident are you that this applicant will repay the loan?”.

$$\begin{aligned} Repay_{pij} &= X_{pij}\beta_{10} + \beta_{0p} + \epsilon_{pij} \\ \beta_{0p} &= \upsilon_{00} + \xi_{0p} \end{aligned} \tag{2.15}$$

Because we did not observe each participant’s implicitly determined expectation of the applicant i’s quality, conditional on score, we used this measure as a proxy. Our reasoning is that if there is discrimination against L group applicants, we should observe lower average *Repay* values for L group applicants than for equally qualified H group applicants.

Tests of Result 2.1 and Proposition 2.1 use the same design matrix, X_{pij} which contains dummy variables for $L{-}Group$, study conditions ($Group$, $QualityLowVar$, $ScoreLowVar$), and their interactions. The design matrix also contains continuous-measure variables that control for applicant quality ($Quality$), score ($Score$), and their interaction. The final two variables included in the design matrix control for time and learning effects ($Round$, $Round^2$). For Propositions 2.2 and 2.3, we compared Study 1 (static inequality) and Study 2 (dynamic inequality) responses of participants from the study condition where group information is available. We modified the design matrix used for Result 2.1 and Proposition 2.1 in the following way: we dropped the $Group$ dummy, added a dummy for $Dynamic$ (indicating Study 2), and interacted $Quality$ with $Dynamic$, $L{-}Group$, $Round$, and $Round^2$ for Proposition 2.1. All continuous variables were first standardized before estimation; therefore, effects associated with these variables should be interpreted in terms of the impact of one standard deviation away from the mean. We used an OLS model to test Result 2.2. We regressed participant p’s total game earnings on $Group_p$.

We included $ScoreLowVar$, $QualityLowVar$, and their interaction to control for study conditions. A positive coefficient on $Group$ would indicate that utilizing group information in addition to individual information is more profitable than just individual information in making service decisions. All models were estimated with standard maximum likelihood methods using the panel data set of participant responses from the two studies. We now turn to discussing results in the next section.

2.4.2 Empirical Results

Figure 2.3 (see after Reference) visually summarizes some of the main results from both studies. The graphs display quality levels of applicants (x-axis) vs. loan offer rates (y-axis). Loan offer rates are the mean percentage of applicants, conditional on quality, that were offered loans by participants across all study conditions. The left column of graphs displays Study 1 results and the right column displays Study 2 results. The top row of graphs results are under the study condition where group membership information is present (Group-Aware) condition. The bottom row results represent the group information is absent (Group-Blind) condition. In both graphs, there is a black vertical line labeled Q^{min} which, unknown to participants, is the minimum applicant quality level needed in order for the participant to earn profit on a loan. All applicant quality values to the right of Q_{min} in the graph are profitable, and those to the left are unprofitable. The curves in each graph with square markers represent loan offer rates for H group applicants while those with

Figure 2.3: Mean Loan Offer Rates

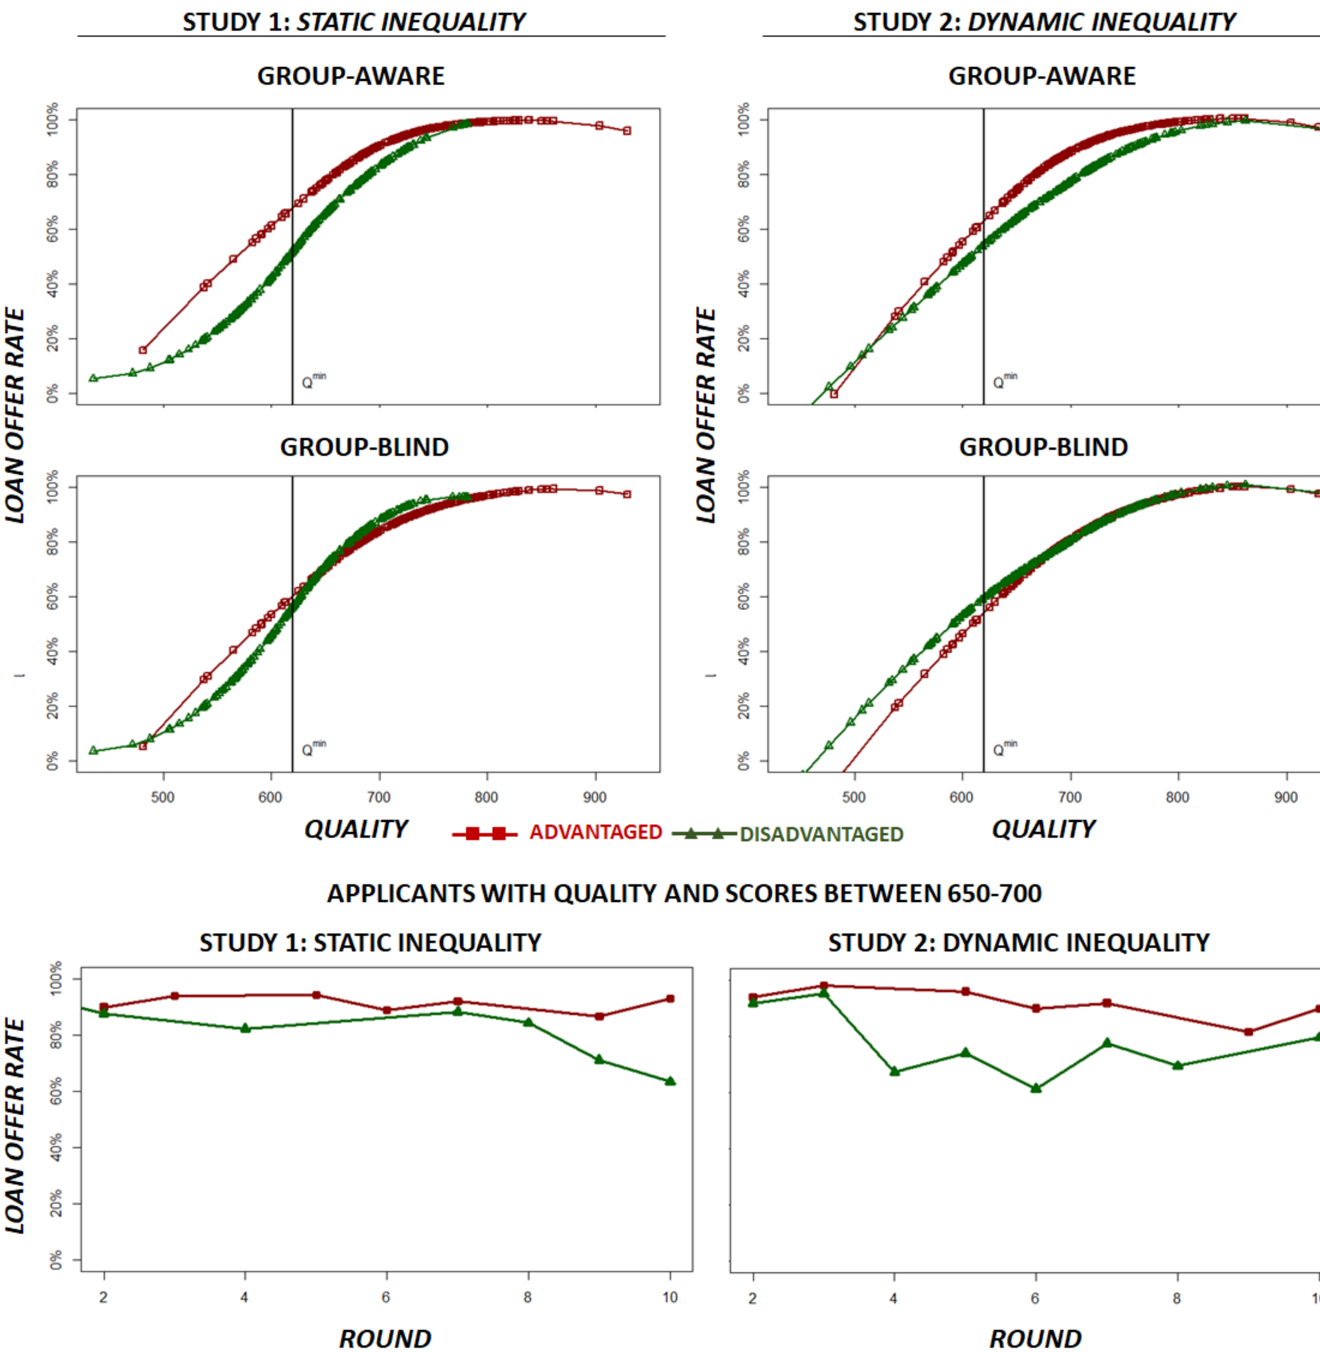

triangle markers represents L group applicants. Because the raw data are visually noisy, we smooth the data with Gaussian LOESS curves (Cleveland and Devlin, 1988; Rust and Bornman, 1982). We use quadratic polynomials with smoothing parameters automatically selected to minimize a bias-corrected AIC (Hurvich et al., 1998). These parameters range from 0.941 to 0.946 and are robust to different degrees of polynomial and to an alternative method of optimized parameter selection (generalized cross-validation). Note from the graphs that in general, loan offer rates for H group applicants dominate those for L group applicants. Although our definition of discrimination is conditional on quality *and* credit score, the gaps between the H group and L group curves give an indication of the magnitude of discrimination between equally qualified H group and L group applicants. One can see that the gaps are greater in the Group-Aware graphs than in the Group-Blind graphs.

Table 2.1 (see after Reference) summarizes the key results of the two studies.

Study 1 and Study 2 results indicate that participants in the Group-Aware condition under both static and dynamic inequality scenarios exact a higher minimum score criterion for L group than equally qualified H group applicants. Participants in the Group-Aware condition were less likely to offer a loan to an L group candidate than to an equally qualified H group applicant (Study 1: -1.013, $p < .001$; Study 2: -1.837, $p < .001$). This finding is consistent with Result 2.1. Studies 1 and 2 are also consistent with Proposition 2.1. When group information is absent, we found no statistical difference in the repayment expectations of equally qualified H group and L group candidates. However, when group information was present and measurement error was low, we find that participants had an increased *Repay* eval-

Table 2.1: Summary of Empirical Results

	Proposition/Result	Dependent Variable	Coefficient of Interest	Study 1 Effects: Static Inequality	Study 2 Effects: Dynamic Inequality	Supported?
1	The service provider's minimum score criterion will be higher for the L group than H group, even for consumers of equal quality.	Pr(Loan)	L-Group*Group	-1.013 p <.001	-1.837 p <.001	Yes
2.1	Discrimination decreases when the error in measuring consumer quality decreases.	Repay	L-Group*Group *ScoreLowVar	2.982 p <.001	3.358 p <.001	Yes
2.2	Discrimination increases when variance in consumer quality decreases.	Repay	L-Group*Group *QualityLowVar	-6.383 p <.001	-1.499 p = .051	Yes
3	The service provider's average profit per transaction (short-term profit) is greater when he uses a Group-Aware policy than a Group-Blind policy.	Profit	Group	.152 p is n.s.	.092 p is n.s.	No
4	If the L group improves in mean quality over time, there are conditions where the group's minimum score criterion rises over time because of the group's improving quality.	Pr(Loan)	L-Group *Dynamic*Round	N/A	-.729 p <.05	Yes
5	An L group's improvement in mean quality over time, will result in some members experiencing decreasing discrimination while others will see increasing discrimination.	Repay	L-Group*Dynamic *Round*Quality	N/A	27.210 p <.05	Yes

uation of L group applicants; this implies a smaller difference in evaluation between equally qualified L group and H group applicants (Study 1: 2.982, $p < .001$; Study 2: 3.358, $p < .001$).

On the other hand, when variance in applicant quality was low, participants had a decreased *Repay* evaluation of L group applicants in Study 1 (-6.383, $p < .001$). In Study 2, the effect of lower variance in applicant quality is not significant when considering all rounds of the study (.565, $p\, is\, n.s.$). However, when considering only the second half of the study, the effect approaches significance and has the expected sign (-1.499, $H_a < 0 \;\; p = .051$). One possible explanation for this result could be that growth in the L group's level of mean quality effectively increases the variation in quality for both low and high variance conditions of the study. Perhaps participants perceived both conditions as high variance. The results suggest that under conditions where growing mean quality (and thus expanding variation in quality) is present, participants needed more time to learn the distinctions between the groups. Furthermore, given that Study 2 has a change in inequality between the applicant groups in each round, it is effectively a 2X2X2X10 study. Statistical power in this study may be an additional factor. This opens up more questions about the long-run effects of time on the perception of variance and its impact on discrimination outcomes. This offers an avenue of interesting future research.

Overall, we do not find statistically significant support for Result 2.2 in either study. Although the coefficient on the *Group* dummy has the expected sign, it is not statistically significant (Study 1: .152; Study 2: .092). We could not say with confidence that, on average, using group information is more profitable than ignoring group information. In fact, we found only one set of conditions where using group information does have a statistically significant profit impact. In Study 1 (where inequality is static), when both variance of quality and variance of measurement error are low, using group information actually *hurts* total participant profits $(-.344,\ p < .05)$. This could be because the credit scores are already very diagnostic in this condition since there is low variation in quality and measurement error. Perhaps group information under these conditions can mislead the participant, especially when its implications are in conflict with those of the credit score.

Study 2 also demonstrates that dynamics in mean quality levels can change minimum score criteria and change the magnitude of discrimination exhibited. Recall that in Study 2, $Q^{min} = 620 < \hat{A}_{L,1} = 640 < \hat{A}_{H,1} = 723$. Given that throughout the study, mean quality levels of both groups exceed Q^{min} and that the L group are growing in quality, we would expect, based on Proposition 2.2, that the minimum score criterion for L group applicants should rise likelihood of loan offers should fall over time when group information is available. We find that the coefficient on the interaction of $Dynamic$, $L - Group$, and $Round$ is consistent with this prediction $(-.729,\ p < .05)$. Despite the fact that the L group grows in quality with each successive round in the game, Study 2 participants were indeed less likely than Study 1 participants to offer loans to equally qualified L group candidates over time.

The bottom two graphs in Figure 2.3 give a visual example of how dynamics in group mean quality can impact the direction of loan decisions in the Group-Aware condition. The graphs display the mean loan offer rates to applicants whose latent quality and credit scores are within the 650 - 700 range. We selected this quality and score range because it contains the largest group of comparably qualified applicants in the study. The left graph shows Study 1 loan offer rates, and the right graph shows the same information from Study 2. H group applicants are represented by the curve marked with squares; L group are marked with triangles. Because their latent quality values exceed $Q^{min} = 620$, all applicants in this group should ideally receive loans. Using the loan offer rates to H group applicants as a point of comparison, one can see that in the static condition, the rate of loan offers to L group applicants gradually decreases during the progress of the study. In contrast, in the dynamic condition the gap between the H group and L group candidates is greater.

Furthermore, we find that increases in discrimination over time are not uniformly experienced across members of the L group. Recall that our study conditions match the conditions of Proposition 2.3 sections 1b and 2b. Under these conditions, we predict that L group members with lower (higher) quality scores would face increasing (decreasing) discrimination in subsequent rounds. We find that the coefficient on the interaction of *Dynamic*, $L - Group$, *Round*, and *Quality* is not significant when considering all rounds of the study ($-.288$). However, when examining the second half of the study (Rounds 6 - 10), the results are significant and support Proposition 2.3 (27.210, $H_a > 0 \;\; p = .042$). As the L group improved in mean quality over time in the study, higher quality L group applicants experienced less discrimination while lower quality applicants experienced more over time. The implications are that although decreasing inequality between groups can lead to reduced discrimination against some L group members, it can also lead to increasing degrees of discrimination against other L group members.

To summarize, the findings from the analytical model, validated by empirical evidence, suggest that when a service provider has access to observable group membership information, the service provider will be less discriminatory against L group consumers from a group with high variance in quality than they will against an L group with low variance in quality. The service provider will also be less discriminatory when the error in measuring consumer quality is low. The findings demonstrate that variability in group members as well as measurement error in detecting quality are each a driving mechanism of discrimination. The findings also show that the service provider will exact a higher minimum score criterion to provide service for L group consumers than for equally qualified H group consumers. Furthermore, the results show that dynamics in mean quality can play a critical role in service discrimination outcomes. If an L group improves its quality over time (which reflects the reality of many U.S. L groups, such as women and minorities), there are certain conditions where the minimum service score criterion can increase despite the group's quality improvement. Moreover, although the L group as a whole may be improving in quality, our results show that not everyone will benefit. While some members will experience decreasing discrimination as a result of the improved quality, others will experience increasing discrimination as a result.

However, the analytical model and empirical studies do not fully address the system complexity over time of discrimination in service and its impact on demand and profits. We next examine the impact of competition, word-of-mouth (WOM), and social factors that can influence variation in customer quality (assimilation, population mix) on the dynamics of service discrimination and on demand and profits. To investigate these, we turn to agent-based modeling. Agent-based modeling (ABM) is a research tool that enables the researcher to simulate the behavior and interactions of autonomous individual agents (people, organizations, etc.) in order to analyze emergent macro phenomena. It is often used to understand the dynamics of collective patterns in a complex system (Delre et al., 2016a; Goldenberg et al., 2001b, 2010; Rand and Rust, 2011). By using both ABM and analytical modeling, we leverage the strengths of each (full parameter space exploration for analytical modeling, modeling of complex interactions for ABM) to answer our research questions more fully than by using one or the other alone (Peres and Van den Bulte, 2014). Using ABM in conjunction with studies can produce new insights through the revelation of macro-level, long-run implications of micro-level observations derived from the studies (Smith and Rand, 2017). ABMs can be used for two different purposes. One purpose is to use an ABM as an extension of an econometric model. In such applications, careful validation of all the input parameters is essential (e.g., see Libai et al., 2013). An alternative use, however, is to use an ABM as an extension of an analytical model in order to show directional results of how variables affect outcomes (Delre et al., 2016b). This reflects our purpose. However, we still strive to use realistic, data-justified values where possible. In that spirit, in the next section we discuss our use of ABM to investigate the long-run implications from our study findings.

2.5 An Agent-Based Model of Service Discrimination

To analyze the dynamics of discrimination in service, we employ a 2^8 full factorial design (256 separate simulations) in the agent-based model (ABM). The ABM models supply and demand for loans in a simulated city. The city contains four competing banks and a population of 200 consumers comprised of people from an H group or L group. Banks and consumers are randomly distributed throughout the geographic area. Based on the distributional assumptions used in the studies, the ABM randomly assigns quality and credit score attributes to consumers. Each bank has one loan officer. Two randomly determined banks have a Group-Aware service policy (a minimum score criterion for each population group) while the others have a Group-Blind service policy (a single minimum score criterion). This allows us to examine competition and its impact on consumer demand and firm profits over time in the ecosystem. Becker (1957) theorized that market forces can ultimately drive out firm discriminatory behavior if non-discriminatory competitors exist. We test the spirit of this theory by including firms in the ABM ecosystem that employ a group-blind minimum score criterion. In each time period in the ABM, a random selection of consumers applies for a loans. These applicants select one and only one bank in any given period based on their utility for the bank (to be elaborated on shortly). Subsequently, each bank loan officer offers loans to applicants with scores exceeding the minimum score criterion determined by bank service policy. Loan officers use historical data of past applicants to update their beliefs about group mean quality levels and to set new minimum score criteria in each period. Each applicant retains a history of loan applications and rejection/acceptance outcomes. Banks cannot observe each applicant's history, but consumers can observe the application history of other consumers in their network.

The ABM uses combinations of high and low values for each of the eight factors. Three of the eight factors come directly from the analytical model and studies: intra-group quality variance, measurement error variance, and degree of inequality ($\hat{\sigma}^2_{qt}$, $\sigma^2_{\varepsilon t}$, and $Inequality_t = \hat{A}_{H,t} - \hat{A}_{L,t}$). We use the same values and decision rules employed in our empirical studies. By doing so, we directly link the empirical study results with the ABM, thereby enabling us to gain insight on the long-run, macro implications of micro-level study observations. Consistent with the empirical studies, we test both static and dynamic inequality conditions over time. In simulations with dynamic inequality, we allow $\hat{A}_L$ to grow at a rate of .16% per period[3] while holding $\hat{A}_H$ fixed.

The remaining five factors are assimilation, population mix, number of applicants, and two dimensions of word-of-mouth (WOM). Assimilation can be thought of as adopting observable characteristics or cultural practices associated with the H group. We expect that greater degrees of assimilation reduce discrimination. Assimilation reduces the chance that an L group member is identifiable as L group because the person possesses attributes of both the H group and L group. For example, a religious minority who attends a bank loan interview dressed in a business suit (characteristic of the H group majority) may experience less discrimination than if

[3]Based on the annual growth rate of average Black wealth relative to Whites from 1967 to 2010 in the U.S. Source–Pew Research Center

he attends in traditional religious garb. We operationalize assimilation in the ABM model by varying the proportion of characteristics (in terms of mean quality) that the L group shares with the H group (0% vs. 50%).

Varying the population mix of the applicant pool allows us to test whether the frequency of exposure to applicants impacts discrimination in service. An increased balance in population mix – a 50/50 split in two populations represents perfect balance – increases the loan officer's exposure to members of both groups. More exposure provides the loan officer with more information. We operationalize population mix in the ABM by varying the percentage of population that is H group (9% vs. 63%)[4]. The lower percentage of 9% represents a less balanced population. We predict that the magnitude of discrimination will be lower when the H group population represents 63% of the population mix. This is because a 63/37 population mix is much closer to a balanced population than a 9/91 split. Discrimination decreases because the loan officer has more information from both groups about consumer quality.

Varying the intensity of demand allows us to test how demand for service impacts service discrimination. We operationalize this by varying the percentage of the city population that applies for a loan in each ABM time period (20% vs. 80%). We posit that a greater frequency of applications would lead to less service discrimination. A greater frequency of applications provides banks with more information. More information should improve variation in quality over time and thus decrease

[4]Based on the percentage of the population that is White in South Africa and U.S respectively. Source: South African National Census of 2011 and 2011 Pew Research Center Report

discrimination. This scenario reflects potential differences between highly trafficked banks (e.g. city banks) versus less trafficked banks (e.g. rural banks), even after controlling for other factors like population mix.

We investigate how the final factor, customer word-of-mouth (WOM), affects demand for services over time. Prior literature has established that WOM can have strong influence on consumer choice (Goldenberg et al., 2001a; Libai et al., 2013; Trusov et al., 2009). Our model assumes loan applicants are utility-maximizing. Utility for bank b has an inverse relationship with distance ($Dist_{ib}$) between applicant i and bank b. It increases with i's assessment of her probability of receiving a loan from the bank. The inclusion of distance as a factor in the utility function is consistent with models in the consumer store choice literature (e.g., Huff, 1964; Rust and Donthu, 1995). We account for additional unobservable factors that influence an applicant's utility with an extreme-valued distributed error term, ε_{ibt}.

WOM about banks is an important factor in each consumer's bank selection. Each consumer in the ABM "talks" to other consumers in her network to find out who has received loans and from which banks. We operationalize WOM through each consumer's ability to access the application history of other consumers in their network. WOM utility that applicant i has for applying to bank b is driven by applicant i's assessed probability that she will be offered a loan from bank b at time t ($P^{WOM}_{ibt} = Pr(Loan_{ibt} \mid \alpha, \mathbf{w_i})$). The probability is equal to the proportion of the applicant's social ties that has received loan offers from bank b weighted by the strength of the social connection between i and each social tie k. Consistent with prior research, strong ties have a greater probability of affecting an individual's choice than weak ties (Brown and Reingen, 1987). The strength of the social connection is measured as the inverse of the distance (Soc_{ik}) between i and k in the simulated city. WOM is also weighted by whether the source of WOM is an in-group vs. out-group member. For example, if i is a member of the squares group in the ABM, then i considers other squares as in-group sources of WOM and triangles as out-group sources. Extant literature has shown that consumers give consideration to in-group versus out-group sources of WOM (Lam et al., 2009; Podoshen, 2006; Uslu et al., 2013).

We vary $\alpha \geq 1$, the weight that consumers place on WOM received from in-group relative to out-group sources, with input values of 1 vs 3 (based on Brown and Reingen (1987); Podoshen (2006); Zhao and Xie (2011) findings). When $\alpha = 1$, applicant i equally weights in-group and out-sources of WOM. An $\alpha > 1$ implies that i places greater weight on WOM from other in-group ties. We also vary β, the weight that consumers place on WOM about bank b relative to the weight placed on the distance to the bank $Dist_{ib}$, with values 2 vs. 20 (based on Trusov et al. (2009) findings). The utility that i has for applying to bank b at time t is as follows:

$$\begin{aligned}
&U_{ibt} = \beta P_{ibt}^{WOM} - Dist_{ib} + \varepsilon_{it} \text{ , where} \\
&P_{ibt}^{WOM} = \frac{\sum_k w_{ik} \ \mathbf{1}(\text{if } b \text{ has ever offered a loan to } k \text{ as of time } t)}{\sum_k w_{ik}} \qquad (2.16) \\
&w_{ik} = \frac{1 + \alpha \ \mathbf{1}(i, k \in j)}{Soc_{ik}}
\end{aligned}$$

Each replication of the bank-applicant ecosystem runs for 300 time periods. Developed in the NetLogo programming language (Wilensky, 1999), the ABM generated over 15.7 million records of data.

2.5.1 ABM Analysis and Results

The results we now share provide additional insight into the dynamics of discrimination in service. Similar to the results from the empirical studies, note that the gap between the H group and L group Group-Aware bank loan offer rates is larger than the Group-Blind gap. Furthermore, the direction and significance of ABM effects are consistent with study results. Consistent with Result 2.1, for example, the Group-Aware banks in the ABM are significantly more likely to offer loans to H group applicants than their L group counterparts (Static Mean Quality: $-1.616, \quad p < .001$; Dynamic Mean Quality: $-.801, \quad p < .001$). Consistent with Proposition 2.1, decreases in measurement error decreases service discrimination (Static Mean Quality: $-24.118, \quad p < .001$; Dynamic Mean Quality: $-14.720, \quad p < .001$)[5]. Lower intra-group variance in quality increases the magnitude of discrimination (Static Mean Quality: $13.774, \quad p < .001$; Dynamic Mean Quality: $16.706, \quad p < .001$). These results provide added confidence that the ABM is appropriately simulating the micro-results from the studies.

Consistent with our prediction, the ABM results suggest that increases in the proportion of the population that is H group (moving from an imbalanced to a balanced, integrated society) decreases discrimination (Static Mean Quality: $-1.982, \quad p < .001$; Dynamic Mean Quality: $-9.214, \quad p < .001$). Recall that discrimination is measured as a difference in expected quality, conditional on two

[5]The dependent variable is an exact measure of discrimination based on Equation 2.11. Data has been 1% trimmed to reduce the effect of extreme outliers of discrimination values.

consumers from two groups having the same quality and score. A greater percentage of the population applying for loans increases discrimination (Static Mean Quality: 19.433, $p < .001$; Dynamic Mean Quality: 13.377, $p < .001$). Decreased assimilation also has the significant effect of increasing discrimination (Static Mean Quality: 46.990, $p < .001$; Dynamic Mean Quality: 33.980, $p < .001$). Recall that the degree of assimilation relates to the proportion of characteristics, and thus mean quality level, that the L group shares with the H group. The ABM results support the expectation that the more assimilated the L group is, the less the group is discriminated against in receiving service.

We find that WOM and competition can drive loss of applicant market share and long-term profits. On average, Group-Blind banks have a significantly greater share of all applicants in the market (Static Mean Quality: 52.4% vs. 47.6%, $p < .001$; Dynamic Mean Quality: 52.8% vs. 47.2%, $p < .001$). WOM also can have a large impact on long-term profits. We regressed cumulative profits on $Group$,α (weight placed on in-group sources of WOM), β (weight placed on WOM in general in the applicant's utility function), and their interactions. We also included controls for other ABM simulation factors ($ScoreLowVar$, $QualityLowVar$, assimilation, population mix, number of applicants). Consistent with findings from prior WOM literature (Libai et al., 2013; Trusov et al., 2009), we find that WOM in general (β) has a positive impact on long-term profits (Static Mean Quality: \$2,333.27, $p < .001$; Dynamic Mean Quality: \$3,537.74, $p < .001$). However, the interaction of WOM parameters with the $Group$ dummy reveals that the greater the weight consumers place on WOM in general, the more negative its impact on the long-term profits of Group-Aware banks relative to Group-Blind banks (Static Mean Quality: \$-4,487.35, $p < .001$; Static Mean Quality: \$-6,431.03, $p < .001$). Regarding the weight placed on in-group sources of WOM (α), we find mixed statistical support of its impact on profits. Overall, the weight on in-group sourced WOM has a directionally positive but not statistically significant impact on long-term profits (Static Mean Quality: \$1,460.26, $p = .170$; Dynamic Mean Quality: \$2,241.08, $p < .091$). However, its effect on Group-Aware banks' long-term profits is negative and statistically significant (Static Mean Quality: \$-2,597.01, $p = .085$; Dynamic Mean Quality: \$-4,176.76, $p = .026$)

Comparing average Group-Aware and Group-Blind banks' short-term profits across all ABM conditions, we find that the Group-Aware banks have, on average, higher profits per loan under static mean quality conditions (Static Mean Quality: \$72.83 Group-Aware vs. \$69.05 Group-Blind, $p < .001$). This is consistent with our hypothesis in Result 2.2 which suggests that discrimination is profitable in the short-run. However, under dynamic mean quality conditions, we find directional but not statistically significant support (Dynamic Mean Quality: \$100.54 Group-Aware vs. \$100.20 Group-Blind, $p = .83$). This is likely because the L group grows in mean quality throughout the simulation to eventually equal the H group population by the end of the simulation. Note that under both static and dynamic mean quality conditions, the difference between the Group-Aware and Group-Blind policies in average profit per period is small. This may provide some indication as to why we were unable to find any statistical difference between participant earnings in the Group-Aware vs. Group-Blind conditions in the empirical studies. The studies have far less statistical power than the ABM.

However, when we compare average Group-Aware and Group-Blind banks' long-term profits across all ABM conditions, we find a reversal. Figure 2.4 (see after Reference) shows average cumulative profits of each type of bank across ABM conditions. On average, Group-Blind banks have sizably greater cumulative prof-

Figure 2.4: Long-Term Profits: Group-Blind vs. Group-Aware Banks

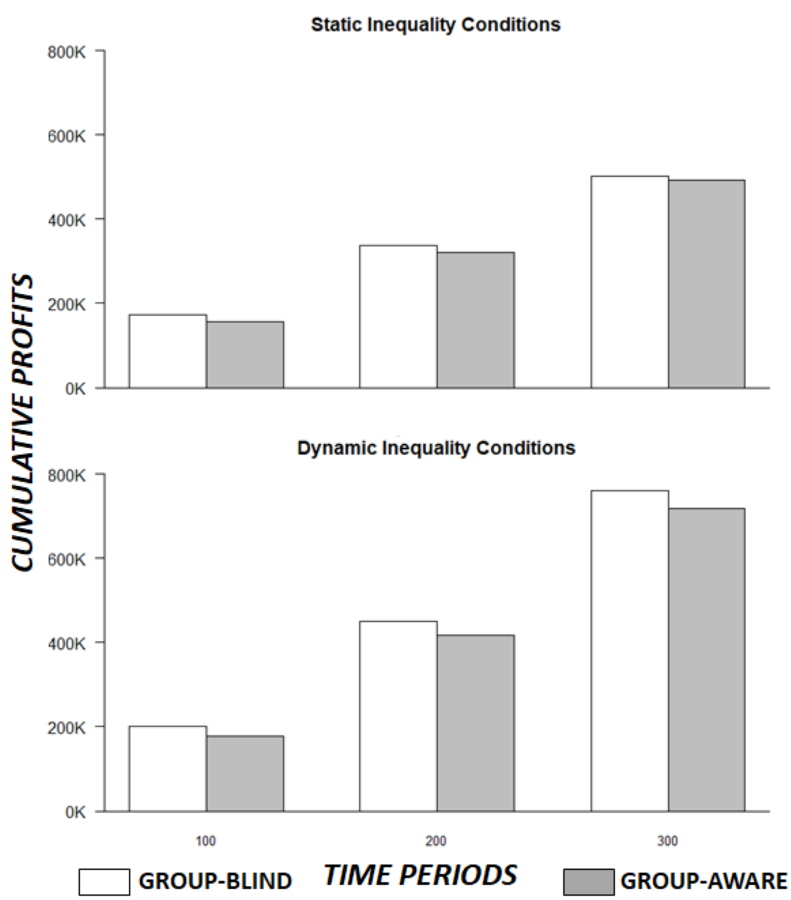

its than Group-Aware banks (Static Mean Quality: \$255,437.73 Group-Blind vs. \$240,966.49 Group-Aware, $p < .001$; Dynamic Mean Quality: \$339,956.23 Group-Blind vs. \$313,239.71 Group-Aware, $p < .001$). By regressing cumulative profits on *Group*, *time*, $time^2$, and their interactions, we find that while the main effect on *Group* (representing Group-Aware banks) is negative but not significant, its interaction with *time* indicates that Group-Aware bank profits substantially erode over time (Static Mean Quality: \$-134.73, $p < .05$; Dynamic Mean Quality: \$-196.93, $p < .01$). In the long-run, myopically profitable, rationally-based discrimination does not pay.

2.6 Discussion

2.6.1 Summary

Our study shows how discrimination in service can emerge from seemingly-rational, non-prejudiced decision-making. We define discrimination in service as different service treatment of equally qualified consumers who differ only in group membership. We distinguish discrimination from prejudice in that prejudice, stereotypes, and racism focus on internally-held attitudes, beliefs, and ideologies. In contrast, discrimination, consistent with sociological literature, is independent of internally-held attitudes. Discrimination concerns decision outcomes that exhibit unequal treatment of people based on the category to which they belong; discrimination is not necessarily driven by internally-held attitudes such as prejudice or bigotry (Pager and Shepherd, 2008; Quillian, 2006).

Although many associate discrimination with race, ethnicity, and gender, our theory and findings should equally apply to many more contexts beyond these categories. They apply to any service scenario where the service provider 1) can segment consumers into groups based on some observable attribute; and 2) the service provider uses group membership as well as individual information about the consumer to make a decision about the provision of service to the consumer. For example, consider how our theory applies to the scenario of the auto salesman who must decide whether to spend his next hour showing Mercedes-Benz E-Class Cabriolets to an 18-year old man versus a 65-year old man waiting in the dealership lobby. Or perhaps the salesman's decision is about a 65-year old in a garbage man's uniform versus a 65-year old man in a business suit. Service decisions such as these, in isolation, may seem to have little impact on firm profits. But the macro social patterns that can emerge from service decisions that rely on group information can produce discriminatory outcomes with negative long-term profit implications.

We illustrate our theory with an example context of bank lending to applicants from either an affluent (H group) or working-class (L group) part of town. Our study demonstrates that discrimination in service can be profitable in the short-run, yet unprofitable in the long-run in competitive markets. This is especially true if consumer word-of-mouth is extensive, as is increasingly the case with modern social media. In our agent-based model, we find that service providers using a Group-Blind service policy that ignores group membership information about consumers have greater total profits over time than those with a Group-Aware service policy that uses group membership information in addition to individual attributes in service decision-making.

2.6.2 Contribution

Our research provides the following theoretical contributions to the literature: First, we examine the critical role that variance plays in the emergence and persistence of service discrimination. Our research shows that service discrimination can arise from low intra-group variation in consumer quality and high measurement error of customer quality. Second, our findings demonstrate that temporal changes in group mean quality level can potentially improve or exacerbate service discrimination. We found conditions where an L group can experience increasing discrimination despite its improving mean quality levels over time. This is of concern because historically L groups have been improving in quality over time in the U.S. Third, we show that if word-of-mouth is influential and if competition can provide outside options to consumers, then a Group-Aware service policy using a minimum score criterion can be more profitable in the short-run, but less profitable in the long run compared to a Group-blind service policy. This matters because a myopic firm can be led down a damaging path by short-term profitability when using group information in its service decisions. Our research also provides a methodological contribution. To the best of our knowledge, we are the first to employ a joint analytical/agent-based model to provide unique insights into the emergent macro effects over time of micro-behavior derived from an analytical model.

2.6.3 Managerial, Consumer, and Public Policy Implications

These findings have important managerial implications. First, we recommend that firms who use a Group-Aware policy in decision-making switch to a Group-Blind policy. The firm should consider the long-term benefits of switching to a Group-Blind service policy that does not use group membership information. We have shown that employing a Group-Blind service policy can provide a strong competitive advantage. It initially seems that a Group-Aware service policy should be more profitable because such a policy provides the service provider an effective device to screen out of risky customers and screen in profitable ones. However, we have shown that such a policy can produce discrimination that erodes profits and market share over time. Because of strong word-of-mouth effects, consumers can learn from each other which firms are unlikely to provide favorable service conditions to them. If services with Group-Blind policies are available as competitive alternatives, L group consumers will switch their preferences for these services over time, and sufficient numbers of H group consumers will patronize Group-Blind services as well. Although discriminatory practices may seem profitable in the short term, they can damage service demand and profits in the long-run.

However, if the firm must persist in using a Group-Aware policy, then we recommend that the firm measure and continually monitor the degree to which there is service discrimination, as well as its impact on profits. Furthermore, we recommend that Group-Aware firms invest in methods of measurement error reduction such as developing advanced methods of measuring consumer quality or more sophisticated predictive models that improve accuracy in predicting quality based on available measures. The Group-Aware firm could also increase its exposure to consumer populations, which could improve information on the mean and variance of group quality. For example, decision-makers could purchase outside data about target markets to supplement its internal data. This could be a way to reduce service discrimination by increasing the decision-maker's exposure to a potentially wider range of consumer quality. This investment should be done at sufficiently frequent intervals with richer predictive models to capture trends in group mean quality levels over time. Another potential solution which may be particularly useful to service providers who implicitly screen customers (e.g. Starbucks, Macy's, Denny's, etc.) is to incorporate in its employee training programs methods and materials that deliberately increase perceived variability of members of different consumer groups. For example, Brauer and Er-rafiy (2011) show that exposing study participants to posters, pictures, articles, and video that highlight the heterogeneity of members of Middle-Eastern and Chinese groups consistently reduced participant discrimination against the each of the groups. By doing so, a firm can put itself on the path to reducing discrimination in service and increasing its profits over time.

These findings also have consumer implications. Our findings imply that consumers seeking less discriminatory experiences in service would do well to seek out services that are, by nature, Group-Blind. For example, many e-commerce sites are more akin to Group-Blind service providers since they have either no access or far less access to group membership information than their bricks-and-mortar counterparts (e.g., think of buying shoes on DSW.com versus walking into a DSW store). Another consumer implication directly results from the knowledge that Group-Aware services are likely to have different service criteria for groups that differ in mean quality. With this knowledge, if a consumer must seek service from a Group-Aware service provider, he or she would do best by masking or omitting information on group membership. Alternatively, the consumer could seek the provider that has the most favorable minimum score criterion for his or her group. The consumer could also improve her outcome by acquiring attributes of the H group (assimilation) when seeking service. For example, the man seeking to buy a Mercedes-Benz at an auto dealership may have a better service experience by wearing a business suit, regardless of his age or occupation.

Our research has public policy implications as well. Currently in the U.S., there are laws, such as the Civil Rights Act of 1964, the Equal Credit Opportunity Act, and the Fair Housing Act, that strive to protect consumers from discrimination in service. However, the task of identifying and proving existence of discrimination in support of enforcing these laws has been a difficult and controversial one. For example, the U.S. Senate recently voted to strike down a rule designed to curb racial discrimination in auto financing. Striking down the rule would provide auto lenders the right to use different score cutoffs for different groups (Merle, 2018). This is a public policy debate which our research addresses precisely. One of the reasons given for repealing the rule is the controversy surrounding how the Consumer Financial Protection Bureau determines whether discrimination exists in the first place (Hayashi, 2018). Our research theory, definition of discrimination, and our findings can provide a framework for developing analytical tools to detect and measure discrimination. Furthermore, the same framework could be the basis of measurement in litigation cases of consumer discrimination.

2.6.4 Limitations and Opportunities for Future Research

There are limitations to this research which suggest many ways that researchers can broaden our knowledge on this topic. For example, our theory assumes customers are members of only one population group. In reality, a consumer can be a member of multiple groups, some of which may be H group while others may not (e.g., a wealthy entrepreneur who has no high-school or college degree). It would be interesting to explore the boundaries of our theory under conditions where consumers may have two or more group memberships with varying levels of mean quality. A second limitation is that we assume in our theory that firms continuously update beliefs using all historical information available about customers who have sought their service. Although we have found qualitative support in our interviews that this can happen in loan services, this may not be true in all service contexts. A promising avenue for future research is investigating how varying the frequency of updates and varying the historical window of data about consumers can affect service provider beliefs. A third limitation of our research is that we assume that the distribution of consumer quality is normally distributed. Although this is a generally reasonable approximation, it would be interesting to explore the effects of other distributional assumptions on service discrimination and on profits. A great deal of work is still needed to fully understand the nature and boundaries of service discrimination, but we believe that the theoretical framework created here serves as a launching point to exploring these and many more questions about the effects of discrimination in service.

2.6.5 Conclusion

We had three goals at the outset of the research discussed in this paper: 1) to uncover the mechanism by which service discrimination can emerge from seemingly rational service policy; 2) to investigate how service discrimination interacts with competition and consumer word-of-mouth to affect profits; 3) to help firms avoid losing profits due to discrimination. We did so by developing a theoretical model that illuminates the critical roles that variation in consumer quality and measurement error in detecting quality play in the emergence and magnitude of discrimination in service. Our theoretical model also demonstrated that changes in group mean quality over time can erode or improve the magnitude of discrimination over time, even if the inequality gap is decreasing. We validated our theoretical model with empirical evidence in two studies. The evidence supported our theory that large variation in consumer quality reduces service discrimination while large measurement error increases service discrimination. Furthermore, the empirical evidence demonstrates that under certain conditions, decreasing inequality between groups can actually *increase* service discrimination. With our agent-based model, we showed the long-term macro effects on profits when firm competition and consumer word-of-mouth embedded in a complex system are taken into consideration. We found that although Group-Blind service providers, who do not use consumer group membership information in its service decisions, are less profitable than their Group-Aware competitors in the short-run, Group-Blind service providers are more profitable in the long-run. This is because consumer word-of-mouth drives consumers to select the most service-friendly alternatives among competitive options.

We provide managerial recommendations on reducing service discrimination's profit-damaging effects. This research emphasizes the long-term benefits of switching to a service policy that does not use group identity information. However, for firms that must persist in using group identity information, this research has additional recommendations which include increasing investment in methods of measurement error reduction and increasing exposure to consumers of different populations. By doing so, a firm could reduce service discrimination while improving its long-term profits and societal well-being.

Appendix A: Competition for Reputation: Category Model Results

Table A.1: Estimates from HB Probit with Gaussian Copulas

	Movies		Music		TV	
INTERCEPT	-11.04	***	-9.17	***	-16.10	**
REPSTOCK	-1.79	***	-1.29	***	-1.13	***
LAWSTOCK	-1.24	***	-1.52	***	-1.25	***
E[SHARE]	-0.51	***	-0.01		-0.15	
E[SHARE] * REPSTOCK	0.33	***	0.52	***	0.20	**
E[SHARE] * LAWSTOCK	0.20	***	0.23	***	0.22	***
REPSTOCK * LAWSTOCK	0.13	***	0.09	*	0.10	*
LAWSTOCK COPULA	2.27	***	2.45	***	2.30	***
REPSTOCK COPULA	7.00	***	5.49	***	7.19	***
E[SHARE] COPULA	19.94	***	24.91	***	32.42	***
Publishers	180		159		110	

	Books		XXX		Applications	
INTERCEPT	-12.74	***	-9.41	***	-39.86	***
REPSTOCK	-1.66	***	-1.55	***	-2.42	***
LAWSTOCK	-0.64	***	-0.52	***	-1.72	***
E[SHARE]	0.62	***	0.52	***	1.06	***
E[SHARE] * REPSTOCK	0.21	***	0.55	***	0.30	*
E[SHARE] * LAWSTOCK	0.16	***	0.19	*	0.18	.
REPSTOCK * LAWSTOCK	0.06		0.09		0.03	
LAWSTOCK COPULA	0.59	*	0.41	.	-0.35	
REPSTOCK COPULA	8.56	***	10.55	***	9.66	***
E[SHARE] COPULA	4.98	***	18.33	***	4.75	***
Publishers	112		74		36	

	Anime		Games		Other	
INTERCEPT	-14.23	***	-25.17	***	-61.69	***
REPSTOCK	-0.27		-3.21	***	-1.05	
LAWSTOCK	-0.07		-2.31	***	-6.46	***
E[SHARE]	0.40	***	0.78	***	1.30	***
E[SHARE] * REPSTOCK	0.32	**	0.20	.	0.19	
E[SHARE] * LAWSTOCK	0.15		0.25	*	0.02	
REPSTOCK * LAWSTOCK	0.07		0.09		1.10	.
LAWSTOCK COPULA	-0.40		1.40	*	13.63	***
REPSTOCK COPULA	6.95	***	9.55	***	13.62	***
E[SHARE] COPULA	3.46	***	6.15	***	-5.23	***
Publishers	35		34		12	

Note: ·p<0.1; *p<0.05; **p<0.01; ***p<0.001

Appendix B: Competition for Reputation: BitTorrent Lawsuit Process

. The copyright lawsuits referred to in this study are against individuals accused of downloading or uploading copyrighted content without authorization by the copyright owner. The individuals are often initially anonymous to the copyright owner. The firm has only an IP address as an identifier. Hundreds, if not thousands, of defendants are can be named in a single lawsuit. Although piracy lawsuits are filed against violators on sites other than BitTorrent-type sites, we concentrate on BitTorrent lawsuits because BitTorrent is the most commonly cited peer-to-peer (P2P) protocol in our lawsuit data. The following gives an overview of a typical process that takes place after an individual accesses unauthorized content online through BitTorrent means. This process is qualitatively the same for only piracy methods other than BitTorrent (e.g., other P2P protocols, streaming, etc.) where only an IP address as an indicator of who pirated the content in question online.

1. Copyright holders often hire tracking firms to monitor BitTorrent sites for uploading and downloading of unauthorized files containing their proprietary content. The copyright owner (CO) provides to the tracking firm a list of names/titles of content they own. After a period of monitoring unauthorized

access of those titles, the tracking firm provides the CO a record of IP addresses associated with unauthorized access of the content the monitoring process.

2. The CO's legal team uses the record of IP addresses as the basis to file charges of copyright violations. The CO requests a motion of discovery in U.S. federal courts in order to subpoena from internet service providers (ISPs) for data on the true identities and contact information associated with the IP addresses. Sometimes ISPs, in an effort to protect the private information of their clients, will instead opt to send a message on behalf of the CO to the individual associated with the IP address (Comcast, 2017). Figure B.1 displays an example of one such message from Comcast: (Farivar, 2013).

Figure B.1: Example of a IP Violation Message from an ISP

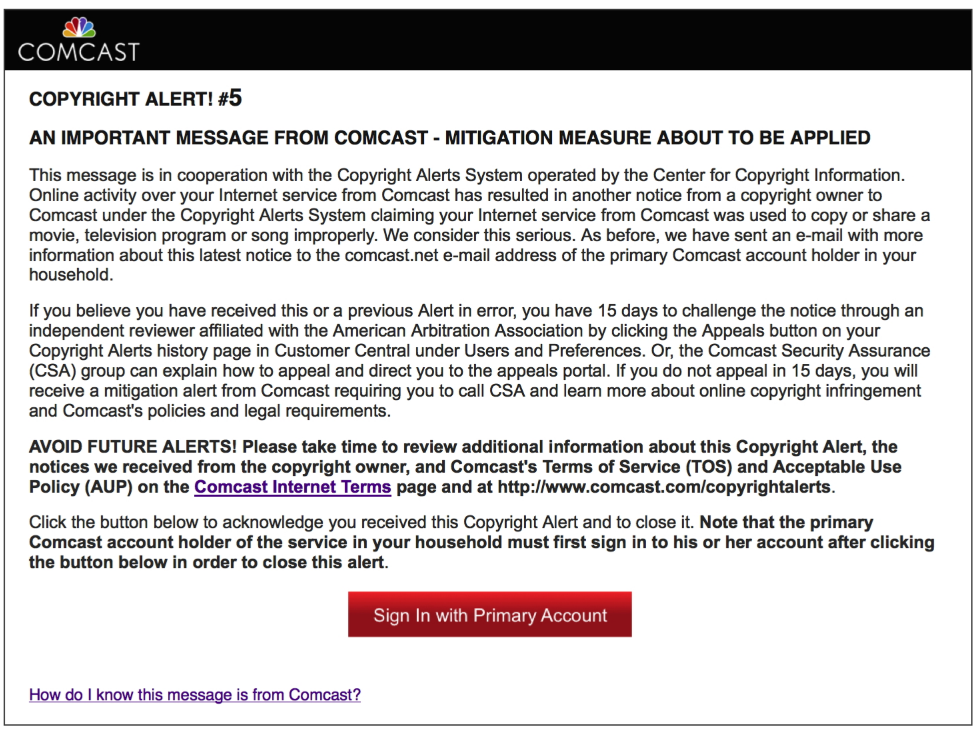

COMCAST

COPYRIGHT ALERT! #5

AN IMPORTANT MESSAGE FROM COMCAST - MITIGATION MEASURE ABOUT TO BE APPLIED

This message is in cooperation with the Copyright Alerts System operated by the Center for Copyright Information. Online activity over your Internet service from Comcast has resulted in another notice from a copyright owner to Comcast under the Copyright Alerts System claiming your Internet service from Comcast was used to copy or share a movie, television program or song improperly. We consider this serious. As before, we have sent an e-mail with more information about this latest notice to the comcast.net e-mail address of the primary Comcast account holder in your household.

If you believe you have received this or a previous Alert in error, you have 15 days to challenge the notice through an independent reviewer affiliated with the American Arbitration Association by clicking the Appeals button on your Copyright Alerts history page in Customer Central under Users and Preferences. Or, the Comcast Security Assurance (CSA) group can explain how to appeal and direct you to the appeals portal. If you do not appeal in 15 days, you will receive a mitigation alert from Comcast requiring you to call CSA and learn more about online copyright infringement and Comcast's policies and legal requirements.

AVOID FUTURE ALERTS! Please take time to review additional information about this Copyright Alert, the notices we received from the copyright owner, and Comcast's Terms of Service (TOS) and Acceptable Use Policy (AUP) on the Comcast Internet Terms page and at http://www.comcast.com/copyrightalerts.

Click the button below to acknowledge you received this Copyright Alert and to close it. **Note that the primary Comcast account holder of the service in your household must first sign in to his or her account after clicking the button below in order to close this alert.**

Sign In with Primary Account

How do I know this message is from Comcast?

If the CO obtains the names and contact information, they send a letter directly to the person associated with the IP address. The letter notifies the person of the copyright infraction and may also provide a recommendation for settlement and resolution with the infringed upon firm in order to avoid potential court proceedings. Figure B.2 displays an example of such a letter from a CO:

Figure B.2: Example of a IP Violation Letter from a Copyright Owner

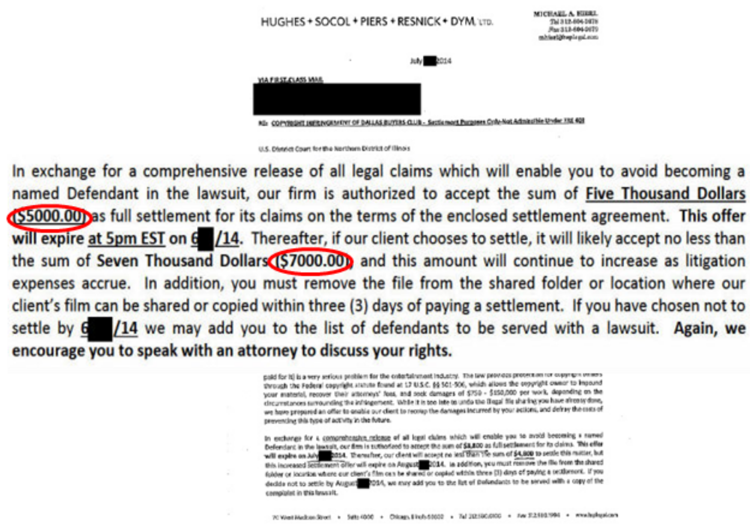
HUGHES + SOCOL + PIERS + RESNICK + DYM, LTD.

July [redacted] 2014

VIA FIRST-CLASS MAIL

[redacted]

U.S. District Court for the Northern District of Illinois

In exchange for a comprehensive release of all legal claims which will enable you to avoid becoming a named Defendant in the lawsuit, our firm is authorized to accept the sum of **Five Thousand Dollars** **($5000.00)** as full settlement for its claims on the terms of the enclosed settlement agreement. **This offer will expire at 5pm EST on 6[redacted]/14.** Thereafter, if our client chooses to settle, it will likely accept no less than the sum of **Seven Thousand Dollars ($7000.00)** and this amount will continue to increase as litigation expenses accrue. In addition, you must remove the file from the shared folder or location where our client's film can be shared or copied within three (3) days of paying a settlement. If you have chosen not to settle by **6[redacted]/14** we may add you to the list of defendants to be served with a lawsuit. **Again, we encourage you to speak with an attorney to discuss your rights.**

3. It is often around this point in the process (filing of initial court paperwork, filing for motion of discovery) that the first wave of news publicity and concomitant word of mouth (WOM) is generated about the pirated content. Sometimes, the firm itself will contribute to the generated publicity by releasing a press release announcing its intent to sue identified individuals. These releases often include statements indicating that the firm intends to search for more violators and add them to the named lawsuit. WOM is evident online in the form of blog posts and comments in BitTorrent and general piracy communities. See Figure B.3 (Gardner, 2011) for some examples of news pieces about

BitTorrent lawsuits.

Figure B.3: Publicity About Online Piracy Lawsuits

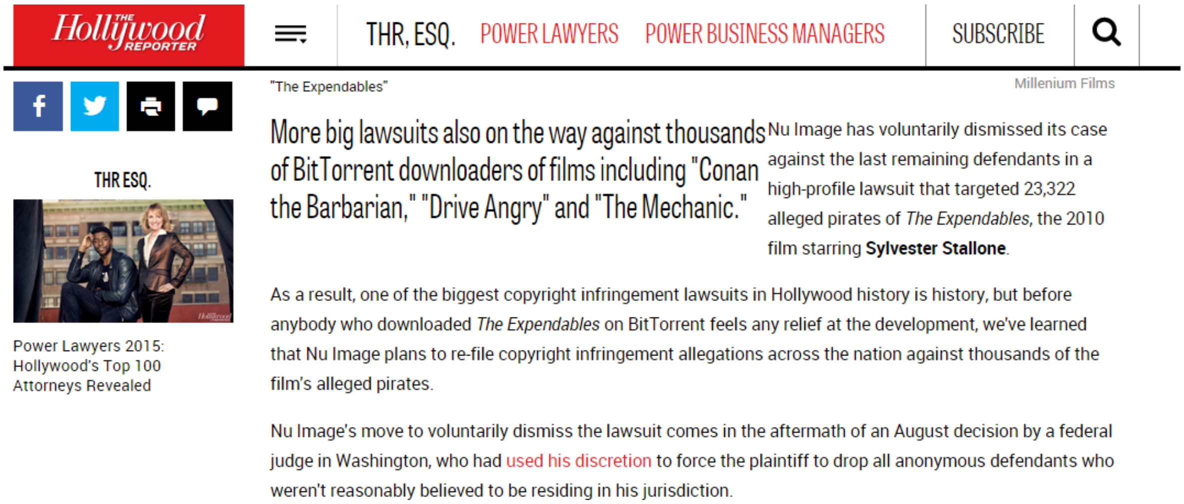

The Hollywood Reporter | THR, ESQ. | POWER LAWYERS | POWER BUSINESS MANAGERS | SUBSCRIBE

THR ESQ.

Power Lawyers 2015: Hollywood's Top 100 Attorneys Revealed

"The Expendables" Millenium Films

More big lawsuits also on the way against thousands of BitTorrent downloaders of films including "Conan the Barbarian," "Drive Angry" and "The Mechanic."

Nu Image has voluntarily dismissed its case against the last remaining defendants in a high-profile lawsuit that targeted 23,322 alleged pirates of *The Expendables*, the 2010 film starring **Sylvester Stallone**.

As a result, one of the biggest copyright infringement lawsuits in Hollywood history is history, but before anybody who downloaded *The Expendables* on BitTorrent feels any relief at the development, we've learned that Nu Image plans to re-file copyright infringement allegations across the nation against thousands of the film's alleged pirates.

Nu Image's move to voluntarily dismiss the lawsuit comes in the aftermath of an August decision by a federal judge in Washington, who had used his discretion to force the plaintiff to drop all anonymous defendants who weren't reasonably believed to be residing in his jurisdiction.

4. Once the firm had received information of the identity associated with the IP address, the firm contacts the individual with a threat letter. Threat letters usually have a similar form and tone, regardless of the attorney offices that issue them. They generally place heavy emphasis on the potential severity of the threat to the recipient. First, they usually inform the individual that he is being sued and outlines the details of the infraction of which the person is being accused. The letters then point out the maximum penalty possible under U.S. law associated with such an infraction. The letters then highlight historical examples of individuals who were successfully sued by firms for similar copyright violation cases and had to pay huge penalties (on the order of hundreds of thousands of dollars). It furthermore offers to settle the lawsuit out of court for a fee. Fees typically tend to be anywhere from \$20 to \$7500 per item uploaded. If the recipient fails to comply by a given date, the settle-

ment fee required to settle increases until the final date given, at which point the firm will pursue litigation to the fullest extent of the law in court, which could include incarceration. Sometimes, the threat letters use shame tactics in addition to monetary or incarceration penalties to heighten the severity of the potential punishment. This is highly dependent on the nature of the content whose copyright was violated. Often times, the content in question is socially embarrassing in nature (e.g. pornography). These letters threaten to make public not only the individuals name and address, but also the titles of the content that he allegedly pirated if he does not pay the fee by a certain date. Threat letters present the second opportunity where publicity is generated about the lawsuit campaign. WOM, and sometimes publicity, are often generated about threat letters sent to individuals. Online WOM is again evident in online social arenas where online pirates participate. For example, the following figure displays an individual posting a copy of his letter with a request for comment or help on how to address it, and whether its threat is credible in the first place (see Figure B.4 (anonymous, 2013)for an example).

Figure B.4: How Online Pirates Learn About Copyright Lawsuits

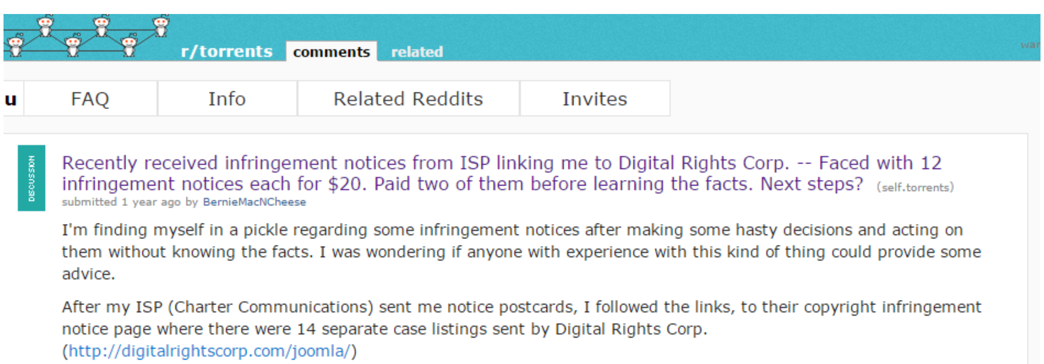

5. A common form of publicity about this stage appears in online blogs that discuss specific threat letters sent to an anonymous source. At times, there are news items about specific threat letters. Because identification of the individual using IP addresses is uncertain, sometimes the individual sued can be, in some rare cases, a nonsensical selected choice. This has lead to the generation of unflattering publicity to the issuing firm. For example, in 2005, the RIAA sued 83 year old Gertrude Walton, who was accused of pirating over 700 songs (many of which were gangsta rap songs) under the user name "smittenedkitten". Ms. Walton, who did not own a computer, had passed away in 2004 (Bangeman 2005).

6. If the identified individual in the lawsuit fails to settle out of court with the firm, the firm proceeds with full litigation with the hope of being awarded significant compensation for damages in U.S. Federal court. This phase in the process is the third opportunity where publicity and subsequent WOM is generated. In these cases, publicity tends to highlight the enormous compensatory awards successfully sued individuals owed to the litigating firms in question and implications for the out-sized realized cost to the individual of each file of content pirated. For example, much publicity surrounded one of the most famous digital piracy cases in U.S. history. Capitol Records, Inc. v. Thomas-Rasset was the first file-sharing lawsuit by major record labels tried in a U.S. court. Jammie Thomas-Rasset, a Minnesota mother of four,

> lost the 2007 case and was ordered to pay $220,000 for pirating 24 songs. In subsequent appeals which went all the way up to the U.S. Supreme Court, Thomas-Rasset lost the cases and at one point was ordered to pay as high at $1.92 million. Her final case reinstated the original damages of $220,000, which is a little over $9,000 per song.

Over 200,000 BitTorrent users have been sued in the U.S. since early 2010, when the first BitTorrent cases were filed in U.S. Courts. Many more thousands have been sued based on copyright infringement using other peer-to-peer file-sharing protocols (e.g. Napster, eDonkey2000, Kazaa, etc.). However, how much such suits affect the expectations of probability and severity of threat publishers face is ambiguous. While there has been plenty of publicity and WOM concerning the presence and severity of punishment from successfully employed threat letters and copyright lawsuits, there have also been non-trivial amounts of publicity and WOM regarding legal "victories" on behalf of the accused infringers. For example, Nu Image, the producers of the 2010 Sylvester Stallone movie The Expendables, had filed a copyright infringement suit in 2011 against over 23,000 individuals. They later dropped the case against all defendants after a Washington federal judge forced the producers to drop the case against anonymous defendants that could not be proven to live in his jurisdiction. Other cases since then have been dropped for similar reasons, or for reasons associated with violating court procedure (e.g. requirements for filing one lawsuit per defendant). Such publicity has muddied the waters, so to speak, on the likelihood potential copyright infractions, even if caught, will materialize into true realization of punishment. This provides a compelling reason to understand the effectiveness of the lawsuit campaigns in curbing piracy in this context.

Appendix C: Competition for Reputation: Analysis of Potential Ceiling Effects

Table C.1 compares, by content category, the total number of uploads from 2010 to 2013 on the website Kickass Torrents to the number of available titles on Amazon in 2013. The number of available titles are based on an Amazon search for all results in the given content category. There are nine content categories of pirated content on the Kickass Torrents website. Table C.1 displays eight of them (Amazon does not sell content in the ninth: XXX). The source of the Amazon data is the Internet Archive Wayback Machine, where we conducted searches of available titles in a content category on the Amazon site in the month of December, 2013. The latest date with available search results were used to construct this table. The data suggests that total publisher uploads on Kickass Torrents were only a fraction of the available content in each category on Amazon in 2013. Given that Amazon is only one of many potential sources of content, the number of titles on Amazon represent a lower-bound on the number of available titles in each content category produced in 2013. Thus, it is highly unlikely that publisher's ran into resource constraints on available content to upload.

Table C.1: 2013 Comparison of Total Uploads vs Amazon Content

	Uploads	Amazon Titles	% of Amazon
Movies	78,478	1,098,734	7.14%
Music	51,138	5,037,435	1.02%
TV	69,138	269,015	25.70%
Books	40,493	41,644,597	0.10%
Applications	12,579	147,151	8.55%
Anime	37,185	52,685	70.58%
Games	8,864	149,058	5.95%
Other	7,405	35,569	20.82%
Total	305,280	48,420,145	0.63%

Appendix D: Competition for Reputation: Endogeneity - Copulas vs Instrumental Variables

We compare the performance of Park and Gupta (2012)'s Gaussian copula method to the Instrumental Variables/Control Function (IV/CF) method using Lewbel (1997)'s method of constructing instrumental variables. We compare the methods in two ways:

1. By simulating data with known parameters and estimating parameters with both methods. Given that the true parameters are known, we can measure and compare the degree of bias each method produces.

2. By estimating the Publisher's choice model for acquiring reputation. We will compare the estimated parameters to determine where the models agree and disagree in their estimation.

We generate the simulated data using the following DGP used in Case 6 of Park and Gupta (2012). This DGP simulates a process with two endogenous regressors, P_{1i} and P_{2i} and an error term, ξ_i. $P^*_{1i}, P^*_{2i}, \xi^*_i$ are the copula regressors that capture the correlation between their corresponding regressors and the error term. The vector of true parameters is $\alpha = \{-1, 1\}$. The simulated data we generate are 1,000 data sets with 500 data points in each set. The DGP is:

$$Y_i = -1 \cdot P_{1i} + 1 \cdot P_{2i} + \xi_i \tag{D.1}$$

$$\begin{pmatrix} \xi_i^* \\ P_{1i}^* \\ P_{2i}^* \end{pmatrix} \sim \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & 0.1 \\ 0.5 & 0.1 & 1 \end{bmatrix} \right) \tag{D.2}$$

$$\xi_i = \Phi^{-1}_{(0,0.5^2)} \left(\Phi(\xi_i^*) \right) \tag{D.3}$$

$$P_{1i} = 1.5 + 0.2 \cdot \ln \left(-\ln[1 - \Phi(P_{1i}^*)] \right) \tag{D.4}$$

$$P_{2i} = 1.0 + 0.3 \cdot \ln \left(-\ln[1 - \Phi(P_{2i}^*)] \right) \tag{D.5}$$

Note that $\Phi^{-1}_{(0,0.5^2)}$ is the inverse distribution function for the normal distribution with mean 0 and variance 0.5^2. The functions for P_{1i} and P_{2i} are the inverse distribution functions of the extreme value distribution.

Table D.1 shows parameter estimates using Gaussian copulas versus instrumental variables. The columns labeled t_{bias} represent the measure of bias produced by the method. It is a t-statistic that equals the absolute standard errors of the parameter estimate from the true values. Both methods recognize the endogeneity associated with P_{1i} and P_{2i}: the copula regressors as well as the control function regressors are statistically significant. Evidence of bias from the OLS and the IV/Control Function methods lies in their t_{bias} values, which are statistically significant (the t-statistics are greater than 2). However, the t_{bias} of the Copulas method indicates that its parameters are not statistically different from the true parameters of the DGP. Hence, the parameters from the Copulas method are consistent.

Next, we compare estimate of our study's empirical model using the two meth-

Table D.1: Comparison of Endogeneity Methods: Simulation Results

	True Parms	OLS	t_{bias}	Copulas	t_{bias}	IV/Control Function	t_{bias}
P1	-1.000	-1.150^{***}	4.054	-1.054^{***}	0.557	-1.254^{***}	4.379
		(0.037)		(0.097)		(0.058)	
P2	1.000	1.332^{***}	5.929	1.109^{***}	0.690	1.440^{***}	4.731
		(0.056)		(0.158)		(0.093)	
$P1^*$				0.252^{***}			
				(0.030)			
$P2^*$				0.192^{***}			
				(0.059)			
P1 IV-CF						1.132^{***}	
						(0.090)	
P2 IV-CF						0.228^{**}	
						(0.109)	
Observations		500		500		500	

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

ods to address endogeneity. In this case, the DGP is not known, only theoretically anticipated. The results of this analysis are in Table D.2: Both the Copulas and IV/Control Function methods produce estimates that are statistically significant on key parameters. Both indicate that there is evidence of endogeneity of the Repstock and E[Share] variables. However, the Copulas method suggests that Lawstock is also endogenous, while the IV/Control Function method does not. Both methods agree in the direction of the sign on all variables except the two endogenous variables: Reputation and E[Share].

Given that Table D.1 indicates that the IV/Control Function method using Lewbel (1997)'s instruments produces biased estimates, this could be an indication of the issues we discussed in Section 1.6.2 about the use of instrumental variables with our random coefficients model. Furthermore the simulation indicates that the Copulas method produces consistent results. For this reason, we assert that the estimation using the Copulas method presents the best estimation of the true parameters associated with the DGP of our empirical study.

Table D.2: Comparison of Endogeneity Methods: Empirical Results

	Copulas	IV/Control Function
Intercept	-11.042***	-9.55***
	(0.998)	(0.736)
Repstock	-1.79***	1.674***
	(0.150)	(0.137)
Lawstock	-1.24***	-0.644***
	(0.082)	(0.056)
E[Share]	-0.507***	2.087***
	(0.085)	(0.159)
E[Share] · Repstock	0.331***	0.105**
	(0.054)	(0.051)
E[Share] · Lawstock	0.195***	0.445***
	(0.038)	(0.038)
Repstock · Lawstock	0.132***	0.123***
	(0.033)	(0.035)
Lawstock Copula	2.27***	
	(0.159)	
Repstock Copula	6.997***	
	(0.274)	
E[Share] Copula	19.936***	
	(0.937)	
Lawstock IV-CF		-0.053
		(0.046)
Repstock IV-CF		0.488***
		(0.085)
E[Share] IV-CF		0.86***
		(0.116)
Publishers	180	180

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

Appendix E: Discrimination in Service: Proofs

For all proofs, we assume the following: Each loan applicant i is a member of one of two population groups $j \in \{L, H\}$. Initially, the H group has a mean quality level that is greater than the L group ($A_H > A_L$). The groups are initially equal in intra-group variation in quality ($\sigma_q^2 = \sigma_{q_H}^2 = \sigma_{q_L}^2$). We also assume that the bank's ability to measure quality is unaffected by changes in composition of the groups (thus σ_ε^2 is constant across groups and across time).

In order to compare Group-Aware and Group-Blind policies in terms of average per period profits, we first must determine the ordinal relationship of S_{all}^{min}, S_L^{min}, and S_H^{min}. Result 2.1 establishes that the minimum score criteria derived from a Group-Aware policy have the ordinal relationship $S_L^{min} > S_H^{min}$ under most conditions. We their ordinal relationship with respect to S_{all}^{min} this in the following Lemma.

Lemma E.1 *The Group-Aware policy minimum score criterion for the L group is always greater than the minimum score criterion of a Group-Aware policy (i.e., $S_{all}^{min} < S_L^{min}$).*

Proof by contradiction: Let us suppose the contrary, that $S_{all}^{min} > S_L^{min}$. Drawing from Equations (2.8) and (2.9), that implies:

$$S^{min}_{all} > S^{min}_{L}$$

$$Q^{min} + \left(Q^{min} - A_{all}\right)\left(\frac{\sigma^2_{\varepsilon}}{\sigma^2_{q_{all}}}\right) > Q^{min} + \left(Q^{min} - A_{L}\right)\left(\frac{\sigma^2_{\varepsilon}}{\sigma^2_{q}}\right)$$

$$\frac{Q^{min} - A_{all}}{Q^{min} - A_{L}} > \frac{\sigma^2_{q_{all}}}{\sigma^2_{q}} \tag{E.1}$$

Since $(Q^{min} - A_{all}) < (Q^{min} - A_{L})$ and $\sigma^2_{q_{all}} > \sigma^2_{q}$, then

$$\frac{Q^{min} - A_{all}}{Q^{min} - A_{L}} < 1 < \frac{\sigma^2_{q_{all}}}{\sigma^2_{q}}, \text{ which is a contradiction.}$$

$$\therefore S^{min}_{all} < S^{min}_{L} \ \forall\ \sigma^2_q, \sigma^2_{\varepsilon}, Q^{min}, A_j$$

E.0.1 Proof of Result 2.2

We wish establish the conditions where $E(\Pi \mid S^{min}_{j}) \geq E(\Pi \mid S^{min}_{all})$: the average per period profit resulting from a Group-Aware service policy is greater than that of a Group-blind service policy. Based on the equations in (2.10), we can expand this inequality and rearrange terms as follows:

$$E(\Pi \mid S^{min}_{j\in\{L,H\}}) > E(\Pi \mid S^{min}_{all})$$

$$\sum_{j\in\{L,H\}} \int_{S^{min}_{j}}^{\infty} \frac{p_j \ E(Q_{ij} \mid S_{ij}) f_j(S) dS}{2 - F_H(S^{min}_{H}) - F_L(S^{min}_{L})} > \sum_{j\in\{L,H\}} \int_{S^{min}_{all}}^{\infty} \frac{p_j \ E(Q_{i} \mid S_{i}) f_j(S) dS}{2 - F_H(S^{min}_{all}) - F_L(S^{min}_{all})}$$

$$\int_{S^{min}_{H}}^{S^{min}_{all}} \frac{E(Q_{i,H} \mid S_{i,H}) f_H(S) dS}{F_H(S^{min}_{all}) - F_H(S^{min}_{H})} + \int_{S^{min}_{all}}^{\infty} \frac{p_H \ [E(Q_{i,H} \mid S_{i,H}) - E(Q_i \mid S_i)] f_H(S) dS}{2 - F_H(S^{min}_{H}) - F_L(S^{min}_{L})}$$

$$> \int_{S^{min}_{L}}^{\infty} \frac{p_L \ [E(Q_i \mid S_i) - E(Q_{i,L} \mid S_{i,L})] f_L(S) dS}{2 - F_H(S^{min}_{all}) - F_L(S^{min}_{all})} + \int_{S^{min}_{all}}^{S^{min}_{L}} \frac{E(Q_i \mid S_i) f_L(S) dS}{F_L(S^{min}_{L}) - F_L(S^{min}_{all})}$$

$$\therefore E(\Pi \mid S^{min}_{j\in\{L,H\}}) > E(\Pi \mid S^{min}_{all})$$

The following proofs involve dynamics. We first present the following additional assumptions: There are two time periods $t \in \{1, 2\}$, two population groups of consumers $j \in \{L, H\}$, and each loan applicant i is a member of one group and applies in one time period only. At time $t = 1$, the H group has an mean quality level that is greater than the L group ($A_{H,1} > A_{L,1}$), the groups are equal in intra-group variation in quality ($\sigma^2_{q_1} = \sigma^2_{q_{H,1}} = \sigma^2_{q_{L,1}}$). We also assume that the bank's ability to measure quality is unaffected by changes in composition of the groups (thus σ^2_ε is constant across groups and across time).

Let us assume that there are two cohorts of applicants where cohort 1 applies for a loan at time $t = 1$ and cohort 2 applies at $t = 2$. Applicants from both cohorts are members of group j. Let p_{jt} represent the proportion of all applicants from group j that are comprised of cohort 1 applicants. This means that the proportion of all j applicants that are in the first cohort is $p_{j1} \in (0, 1)$ and cohort 2 is $(1 - p_{j2})$. Also assume that the two cohorts are equal in intra-cohort variation in quality ($\sigma^2_{q_{j1}} = \sigma^2_{q_{j2}}$). However, the mean quality level of cohort 2 is g_j times the level of cohort 1 ($A_{j2} = g_j A_{j1}$, where $g_j \in [0, \infty)$).

E.0.2 Proof of Proposition 2.2

If the loan officer uses all available information about group j as of time $t = 2$, then based on Lemma 2.1 and Equation (2.8), the loan officer's minimum score criterion for group j is

$$S^{min}_{jc} = Q^{min} + \left(Q^{min} - A_{jc}(g_j)\right) \left(\frac{\sigma^2_\varepsilon}{\sigma^2_{q_{jc}}(g_j)}\right)$$

To understand the impact of growth of group j's mean quality on the minimum score criterion, we take the derivative of S_{jc}^{min} with respect to g_j.

$$\frac{\partial S_{jc}^{min}}{\partial g_j} = \frac{A_{j1}(1-p_{j1})\sigma_\varepsilon^2}{\sigma_{q_{jc}}^2(g_j)}\left[\frac{2A_{j1}p_{j1}(1-p_{j1})(1-g_j)(Q^{min}-A_{jc}(g_j))}{\sigma_{q_{jc}}^2(g_j)} - 1\right] \tag{E.2}$$

Trivially, $S_{jc}^{min} = S_{j1}^{min}$ when $g_j = 1$. Otherwise, when $g_j \neq 1$, $\frac{\partial S_{jc}^{min}}{\partial g_j}$ has the following behavior, which depends on the relationship of Q^{min} with respect to a threshold value A^*:

$$\frac{\partial S_{jc}^{min}}{\partial g_j}\begin{cases} > 0 & \begin{cases}\text{when } g_j > 1 \text{ and } Q^{min} < A^* \\ \text{when } g_j < 1 \text{ and } Q^{min} > A^*\end{cases} \\ = 0 & \text{when } g_j \neq 1 \text{ and } Q^{min} = A^* \\ < 0 & \text{otherwise} \end{cases} \tag{E.3}$$

where

$$\begin{aligned} A^* &= A_{jc}(g_j) + \frac{\sigma_{q_{jc}}^2(g_j)}{2A_{j1}p_{j1}(1-p_{j1})(1-g_j)} \\ &= A_{j1}\left[p_{j1}(1-g_j) + \frac{1}{2}(1+g_j)\right] + \frac{\sigma_{q_{j1}}^2}{2A_{j1}p_{j1}(1-p_{j1})(1-g_j)} \end{aligned}$$

E.0.3 Proof of Proposition 2.3

We wish to show that a consumer can experience an increasing degree of discrimination over time even if her group's mean quality is improving. We establish this with the following proof. Let consumers i and $-i$ from groups j and $-j$ have constant quality level Q^*. Recall from Definition 2.1 that discrimination is defined as:

$$D_{it} = (\gamma_{H,t} - \gamma_{L,t})Q^* + [(1-\gamma_{H,t})A_{H,t} - (1-\gamma_{L,t})A_{L,t}] \tag{E.4}$$

We define the change in discrimination i experiences over time as

(E.5)

$$\Delta D_i = D_{i2} - D_{i1} = (\Delta\gamma_H - \Delta\gamma_L)\, Q^* + [\Delta(1-\gamma_H)A_H - \Delta(1-\gamma_L)A_L]$$

$$\text{where} \quad \Delta\gamma_j = \gamma_{j2} - \gamma_{j1} \text{ and } \Delta(1-\gamma_j)A_j = [(1-\gamma_{j2})A_{j2} - (1-\gamma_{j1})A_{j1}]$$

If $\gamma_{j2} = \gamma_{j1}$, then all consumers i experience no change in discrimination over time. However, if $\gamma_{j2} \neq \gamma_{j1}$, then the consumer i that experiences no change in discrimination ($\Delta D_i = 0$) has quality

$$Q^* = Q^*_{\Delta D0} = \frac{\Delta(1-\gamma_L)A_L - \Delta(1-\gamma_H)A_H}{\Delta\gamma_H - \Delta\gamma_L} \tag{E.6}$$

Consumers with $Q^* \neq Q^*_{\Delta D0}$ experience changing discrimination under the following conditions:

$$\begin{aligned} \Delta D_i > 0 &\begin{cases} \Delta\gamma_L > \Delta\gamma_H & \text{when } Q^* < Q^*_{\Delta D0} \\ \Delta\gamma_L < \Delta\gamma_H & \text{when } Q^* > Q^*_{\Delta D0} \end{cases} \\ \Delta D_i < 0 &\begin{cases} \Delta\gamma_L > \Delta\gamma_H & \text{when } Q^* > Q^*_{\Delta D0} \\ \Delta\gamma_L < \Delta\gamma_H & \text{when } Q^* < Q^*_{\Delta D0} \end{cases} \end{aligned} \tag{E.7}$$

Hence, different members of the same group j can experience different degrees of discrimination over time as a result of their group's change in mean quality.

Appendix F: Discrimination in Service: Design of Experimental Study

The two experimental studies seek to answers to the following:

- Under what conditions can discrimination in service emerge?
- What are the impact of dynamics in mean group quality on the magnitude of discrimination?

We predict that our findings will support the supposition that this phenomenon arises when mean group quality (reputation) is used as a source of information in the firm's decision-making process about level of service offered to the prospective customer. Specifically, our research investigates and tests the following propositions:

P2.1.1: As the variation of credit scores increases, (implying increasing error in detecting true quality level), discrimination increases.

P2.1.2: As the variation of true quality of individuals increases within each group, discrimination decreases.

P2.2: Growth over time in the L group's mean quality will increase the minimum score criterion for the L group over time.

P2.3: Growth over time in the L group's mean quality will increase the discrimination against the L group over time.

Design of Two Experiments The purpose of our 2 studies is to test for the existence and magnitude of discrimination in service provision to two customer segments, and how mean and variation of group quality moderates it. The context of the study is provision of loans to individuals.

Customers who are members of Group A have, on average, higher mean quality levels than those in Group B. This translates into better loan repayment abilities (higher quality applicants) and higher credit scores relative to group B. Yet, the two groups overlap in these attributes. Thus, it is possible for some Group A customers to have the same loan attributes as some Group B customers. We test whether participants treat loan applicants who are comparable in quality differently when provided group membership information.

Study 1 (Static Advantage):

Description A 2 (no group information provided vs. group information provided) X 2 (low vs. high variation in credit scores, conditional on quality level) X 2 (low vs. high variation in true quality) between subjects design where the mean quality of each group is held constant over time.

Purpose To test P2.1.1 and P2.1.2.

Participants 400 (8 cells X 50 participants per cell)

Methodology

DV 1) Loan given or not (binary) 2) Profits

IV Loan applicant group membership, loan applicant credit score, demographics of the participant

Manipulations

a) Low vs. High variation in Credit Scores, conditional on quality level: Participants in the High (Low) condition will see credit scores generated from a distribution with high (low) variance of credit score, given a quality level (true ability to repay loan). This means that conditional on a group of loan applicants who have the same true ability to repay the loan on average, participants in the High condition will see more variation in the credit scores than those in the Low condition. This higher variation condition can be interpreted as the credit score possessing more error in indicating the true repayment ability of the applicant than scores in the Low condition.

b) Low vs. High variation in Quality Scores, conditional on credit score accuracy level: Participants in the High (Low) condition will see credit scores applicant pools that differ in their true ability to repay the loan. This means that conditional on credit score's accuracy in indicating true loan repayment ability (applicant quality) is held constant, participants in the High condition will see more variation in the true repayment abilities of applicants than those in the Low condition.

c) Group Information Provided vs. No Group Information Provided: Participants in the Group Information Provided condition will see information about attributes of Group A and Group B at the beginning of the game. Those in the No Group Information condition will not see any information about group attributes at the beginning of the game.

Procedure

Each participant will play a simplified lending game, where the participant takes the role of a lender. The participant is guaranteed $4 to participate, no matter the outcome of the game. The participant can increase her compensation investing $3 in a series of loans.

The participant will be presented with a series of applications for a loan to buy a house. The participant will be told that each application has two pieces of information – the applicant's credit score and group membership (A or B). The participant is asked to make a decision about whether or not to lend a fixed amount of money to the applicant. The following shows an example of an applicant profile that a participant might see during the experiment:

After each round of the game is played (in each round, the participant reviews 10 applications), the participant is presented with a table revealing how each applicant who was offered a loan performed. If the loan was repaid, the participant earns $0.03. If the loan defaulted, the participant loses $0.03. The participant plays 10 rounds, after which the total profit made is calculated and paid to the participant. The maximum compensation to the participant is $10, and minimum is $4.

For participants in the Group Information Provided condition, the participant

Figure F.1: Example of Applicant Profile in Experiment

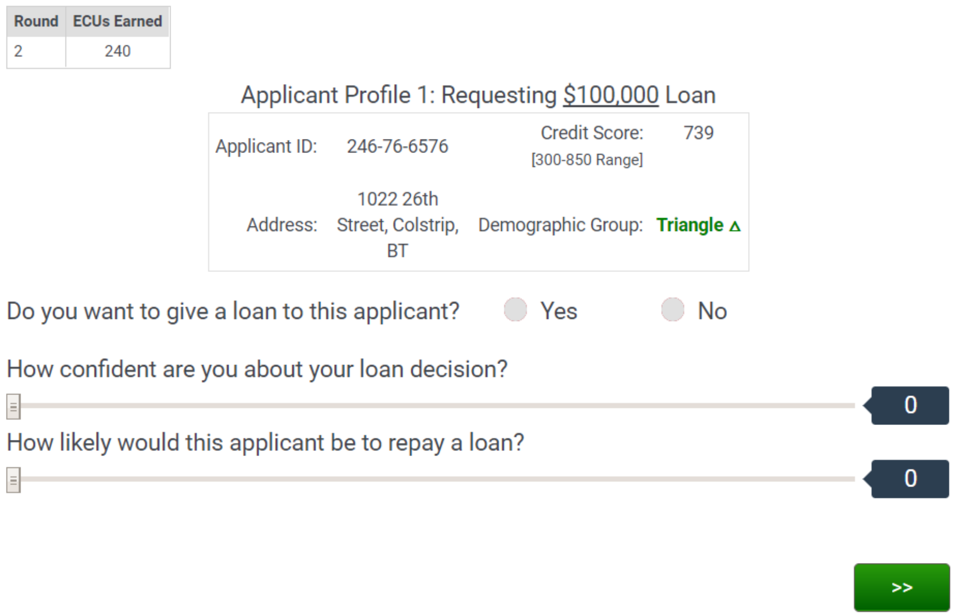

is shown a full page of attributes about the two groups before the game begins (much of which is not relevant to the experiment. This is intentional in order to ensure that the main information of the experiment, the loan historical performance table, is not too salient and leading). In the No Group Information Provided condition, the participants will not receive this information.

All participants will see a table that includes a sample of historical performance of loans given by the bank. The table of historical performance will look something like this:

We intend to carefully construct the information shown in the table so that it is consistent with the distribution of scores and repayment probabilities associated with the true distributions of scores from the high mean quality group A and low mean quality group B.

Figure F.2: Loan Performance Table

Round	ECUs Earned
Practice	0

The practice round has ended. Below are tables showing summary information of loans you gave and of all applications to the bank so far. Reminder: this round does not really count towards ECUs earned.

Applicant That You Gave Loans	Demographic Group	Credit Score (300 – 850 range)	Repaid Loan?	ECUs
1	Square □	795	Yes	30
2	Triangle △	552	Yes	30
3	Triangle △	448	No	-30
4	Square □	560	Yes	30
5	Triangle △	465	No	-30
Total				30

Entire Bank History	Demographic Group	Average Credit Score – All Applicants (300 – 850 range)	Percent Given Loans	Average Credit Score – Loans(300 – 850 range)	Loan Repayment Rate
Average	Square □	678	100%	678	100%
Average	Triangle △	488	100%	488	33%

>>

Study 2: (Dynamic Advantage)

Description A 2 (no group information provided vs. group information provided) X 2 (low vs. high variation in credit scores conditional on quality level) X 2 (low vs. high variation in true quality) between subjects design where the mean quality of group A is held constant while the mean quality of group B increases deterministically over the 10 rounds of play.

Purpose To replicate P2.1.1 and P2.1.2 results from Study 1. To test P2.2 and P2.3.

Participants 400 (8 cells X 50 participants per cell)

Methodology

Same as Study 1. Only differences between Study 1 and Study 2 are as follows:

Profiles of group B members will differ because scores at time t are generated from increasing mean quality levels for group B

The Group B profiles that default on their loans and expected payoffs will change based on the mean quality level at time t.

Manipulations

Same as Study 1.

Proposed Statistical Analysis

1. ANOVA on a) number of loans given and c) loan profitability by participants relative to expected numbers if discrimination did not exist (i.e., ANOVA on the delta between actuals and expected numbers of loans accepted based on true quality of applicants). Tests P2.1.1 and P2.1.2.

2. ANOVA comparing study 1 and study 2 outcomes for each of 4 cells. Tests [P2.2.

3. Probit to analyze contribution of IVs to probability of selecting A/B group applicants. DID regression to analyze contribution of IVs to firm profitability

Appendix G: Discrimination in Service: ABM Rules of Engagement

A Group-Blind bank has a loan strategy that does not use demographic group membership information. Instead, it uses the historical performance of all customers collectively as well as information about the individual applicant in its decision of whether or not to offer a loan.

A Group-Aware bank has a loan strategy that does use demographic group membership information. It uses the historical performance of all customers collectively by group as well as information about the individual applicant, conditional on group membership, in its decision of whether or not to offer a loan.

Customers who are interested in applying for a loan have a strategy where they will choose the bank that maximizes the customer's utility function. Their choice strategy, modeled by a multinomial logit choice function, has a utility that includes the likelihood a bank accepts their application, the distance between the customer's home and the bank, and unobservables captured by an error term. They assess the probability of acceptance through two information sources: advertising and word-of-mouth (WOM). The customer's personal assessment of the probability her application will be accepted depends on the information she gathers via WOM about acceptance rates of each bank. She differentially weights WOM information based on the characteristics of its source (strong vs. weak ties, in-group vs. out-group sources). Literature has demonstrated that people place greater weight on WOM from strong ties versus weak ties (Granovetter 1973, Goldenberg et al., 2001). Literature has also shown that people place greater weight on in-group sources of WOM than out-group sources (Podoshen 2006, Uslu, Durmuş, Taşdemir 2013). If the customer decides to apply for a loan, she selects the option that maximizes her utility. Her choice set is either one of four banks in her world or the outside option (no loan application).

Timeline of events during setup of the world in the ABM simulation (this happens once in simulation)

1. Two Colorblind Banks and two Full Information Banks are randomly placed in geographic locations in the ABM world

2. The banks are endowed with rules for their respective strategies to assess loan applicants (see next list for more detail)

3. A large population of people are randomly distributed throughout geography of world. They are randomly "born" to be either green or red people.

4. Each person is endowed with randomly provided characteristics (e.g., credit score, group membership) as well as rules governing the strategy for applicant's choice selection of a bank to apply for a loan

Timeline of events within one period in the ABM simulation

1. A randomly selected percentage of the population decides to apply for a bank

loan.

2. These applicants select from a choice set of 4 possible banks to apply for the loan based on utility maximization

    (a) Applicant gathers WOM information about each bank in choice set from strong and weak ties. Within each category of ties, applicants also use in-group and out-group WOM information. WOM is regarding the percentage of past ties whose loan was accepted by each bank.

    (b) Applicant computes her own likelihood of her loan being accepted at each bank in choice set based on gathered WOM. Strong ties and in-group WOM gets greater weight than weak ties and out-group WOM.

    (c) Applicant computes her own utility for each bank in choice set, which is based on loan acceptance probability, distance to bank, and error.

    (d) Applicant selects bank based on multinomial logit choice model.

3. The banks review each loan application and forms expectation of applicant quality (interpreted as ability to repay loan) based on its own policy

    (a) Colorblind Banks' policy is to form expectation based on credit score of applicant and historical scores of all past applicants.

    (b) Full Information Banks have policy of forming expectation based on credit score of applicant and historical scores of all past applicants from applicant's group.

4. Banks offer loan to applicant if the expected quality of applicant meets or exceeds banks minimum threshold of quality (Q_{min}). All other applicants are rejected. NOTE: All banks have same minimum threshold.

5. Banks update historical information with applicant's information.

   (a) Colorblind Banks update historical information on scores of all past applicants

      i. Update historical mean of scores S

      ii. Update minimum accepted score S_{min}.

   (b) Full Information Banks update historical information on scores of all past applicants by group membership

      i. Update historical mean of scores S_j by group

      ii. Update minimum accepted score S_{min_j} by group.

6. All applicants who have applied for loan update their own historical information about banks (success/no success at applying for loan at each bank).

   (a) If applicant is rejected by a bank, she updates her probability of loan acceptance by that bank to 0. This remains in effect indefinitely.

   (b) If applicant is accepted by a bank, she updates her probability of loan acceptance by that bank to 1. This remains in effect indefinitely.

   (c) For any bank where the applicant has never applied for a loan, the applicant will update probability of loan acceptance based on gathered WOM information in the next period.

7. Simulation clock proceeds to next period. Entire process starts again.

Bibliography

Addady, M. (2016). 3 More Fast Food Workers Who Refused to Serve Law Enforcement.

Aigner, D. J. and Cain, G. G. (1977). Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review*, 30(2):175.

Amemiya, T. (1973). Regression Analysis when the Dependent Variable Is Truncated Normal. *Econometrica*, 41(6):997.

Anderson, L. and Ostrom, A. L. (2015). Transformative Service Research: Advancing Our Knowledge About Service and Well-Being. *Journal of Service Research*, 18(3):243–249.

Anderson, L., Ostrom, A. L., Bitner, M. J., Brown, S. W., Burkhard, K. A., Goul, M., Smith-Daniels, V., Demirkan, H., and Rabinovich, E. (2010). Improving Well-Being through Transformative Services. *Journal of Service Research*, 13(1):4–36.

anonymous (2013). reddit/torrents - Recently received infringement notices from ISP linking me to Digital Rights Corp. – Faced with 12 infringement notices each for $20. Paid two of them before learning the facts. Next steps?

Antonovics, K. (2006). Statistical discrimination and intergenerational income mobility.

Arrow, K. J. (1973). The theory of discrimination. In *In Discrimination in Labor Markets*, pages 3–33. Princeton University Press.

Bajari, P. and Hortaçsu, A. (2004). Economic Insights from Internet Auctions. *Journal of Economic Literature*, 42(2):457–486.

Bandura, A. (1969). Principles of behavior modification.

Becker, G. S. (1957). *The Economics of Discrimination*. 1 edition.

Becker, G. S. (1968). Crime and punishment: An economic approach. In *The economic dimensions of crime*, pages 13–68. Springer.

Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105(3):921–57.

Berry, S. and Reiss, P. (2007). Chapter 29 Empirical Models of Entry and Market Structure. *Handbook of Industrial Organization*, 3:1845–1886.

Bhattacharjee, S., Gopal, R. D., Lertwachara, K., and Marsden, J. R. (2006). Impact of Legal Threats on Online Music Sharing Activity: An Analysis of Music Industry Legal Actions. *The Journal of Law and Economics*, 49(1):91–114.

Bhutta, N. and Canner, G. B. (2013). Mortgage market conditions and borrower outcomes: Evidence from the 2012 HMDA data and matched HMDA-credit record data. *Federal Reserve Bulletin*, 99(4).

Bjerk, D. (2008). Glass ceilings or sticky floors? Statistical discrimination in a dynamic model of hiring and promotion. *The Economic Journal*, 118(530):961–982.

Blake, Andrew (2016). Kickass Torrents targeted by DOJ, feds say website distributed $1 billion in copyrighted works. *The Washington Times*.

Blume, L. E. (2005). Learning and statistical discrimination. *The American Economic Review*, 95(2):118–121.

Blume, L. E. (2006). The dynamics of statistical discrimination. *The Economic Journal*, 116(515).

Board, S. and Meyer-ter Vehn, M. (2013). Reputation for Quality. *Econometrica*, 81(6):2381–2462.

Bohren, J. A., Imas, A., and Rosenberg, M. (2017). The Dynamics of Discrimination: Theory and Evidence.

Bone, S. A., Christensen, G. L., and Williams, J. D. (2014). Rejected, Shackled, and Alone: The Impact of Systemic Restricted Choice on Minority Consumers' Construction of Self. *Journal of Consumer Research*, 41(2):451–474.

Bone, S. A., Williams, J. D., and Christensen, G. L. (2010). When Consumer Well-Being Meets Small Business Ownership: Transforming Financial Service Systems to Eradicate Disparate Treatment and Discrimination. *Journal of Macromarketing*, 30(4):408–422.

Brauer, M. and Er-rafiy, A. (2011). Increasing perceived variability reduces prejudice and discrimination. *Journal of Experimental Social Psychology*, 47(5):871 – 881.

Braun, M. and Moe, W. W. (2013). Online Display Advertising: Modeling the Effects of Multiple Creatives and Individual Impression Histories. *Marketing Science*, 32(5):753–767.

Broadwater, L. (2012). Wells Fargo agrees to $175 million settlement in pricing discrimination suit. *The Baltimore Sun*.

Brown, A. and Patten, E. (2017). The narrowing, but persistent, gender gap in pay.

Brown, J. J. and Reingen, P. H. (1987). Social ties and word-of-mouth referral behavior. *Journal of Consumer research*, 14(3):350–362.

Burmester, A. B., Becker, J. U., van Heerde, H. J., and Clement, M. (2015). The impact of pre- and post-launch publicity and advertising on new product sales. *International Journal of Research in Marketing*, 32(4):408–417.

Cabral, L. and Hortaçsu, A. (2010). The Dynamics of Seller Reputation: Evidence from Ebay*. *The Journal of Industrial Economics*, 58(1):54–78.

Carroll, A., Houghton, S., Durkin, K., and Hattie, J. A. (2009). *Adolescent reputations and risk: Developmental trajectories to delinquency*. Springer Science & Business Media.

Chevalier, J. A. and Mayzlin, D. (2006). The Effect of Word of Mouth on Sales: Online Book Reviews. *Journal of Marketing Research*, 43(3):345–354.

Chib, S. and Greenberg, E. (1995). Understanding the Metropolis-Hastings Algorithm. *The American Statistician*, 49(4):327–335.

Chung, D. J. (2013). The Dynamic Advertising Effect of Collegiate Athletics. *Marketing Science*, 32(5):679–698.

Cleveland, W. S. and Devlin, S. J. (1988). Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83(403):596–610.

Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 83(5):1220–1240.

Cole, J. R., Cole, S., and deB. Beaver, D. (1974). Social Stratification in Science. *American Journal of Physics*, 42:923–924.

Comcast (2017). Comcast's DMCA Repeat Infringer Policy for Xfinity Internet Service.

Craig, A. C. and Fryer, R. G. (2017). Complementary Bias: A Model of Two-Sided Statistical Discrimination.

Crockett, D., Anderson, L., Bone, S. A., Roy, A., Wang, J. J., and Coble, G. (2011). Immigration, Culture, and Ethnicity in Transformative Consumer Research. *Journal of Public Policy & Marketing*, 30(1):47–54.

Crockett, D., Grier, S. A., and Williams, J. A. (2003). Coping with marketplace discrimination: An exploration of the experiences of black men. *Academy of Marketing Science Review*, 2003:1.

Danaher, B. and Smith, M. D. (2014). Gone in 60 seconds: The impact of the Megaupload shutdown on movie sales. *International Journal of Industrial Organization*, 33:1–8.

Danaher, B., Smith, M. D., Telang, R., and Chen, S. (2014). The Effect of Graduated Response Anti-Piracy Laws on Music Sales: Evidence from an Event Study in France. *J Ind Econ*, 62(3):541–553.

Dellarocas, C., Fan, M., and Wood, C. A. (2004). Self-Interest, Reciprocity, and Participation in Online Reputation Systems. SSRN Scholarly Paper ID 585402, Social Science Research Network, Rochester, NY.

Delre, S. A., Broekhuizen, T. L., and Bijmolt, T. H. (2016a). The effects of shared consumption on product life cycles and advertising effectiveness: The case of the motion picture market. *Journal of Marketing Research*, 53(4):608–627.

Delre, S. A., Panico, C., and Wierenga, B. (2016b). Competitive strategies in the motion picture industry: An abm to study investment decisions. *International Journal of Research in Marketing*, 34(1):69 – 99.

Dickinson-Delaporte, S. J. and Holmes, M. D. (2011). Threat appeal communications: The interplay between health resistance and cognitive appraisal processes. *Journal of Marketing Communications*, 17(2):107–125.

Diekmann, A., Jann, B., Przepiorka, W., and Wehrli, S. (2014). Reputation Formation and the Evolution of Cooperation in Anonymous Online Markets. *Am Sociol Rev*, 79(1):65–85.

DiPrete, T. A. and Eirich, G. M. (2006). Cumulative Advantage as a Mechanism for Inequality: A Review of Theoretical and Empirical Developments. *Annual Review of Sociology*, 32(1):271–297.

Dubé, J.-P., Hitsch, G. J., and Manchanda, P. (2005). An empirical model of advertising dynamics. *Quantitative marketing and economics*, 3(2):107–144.

Eaton David H (2007). The Impact of Reputation Timing and Source on Auction Outcomes. *bejeap*, 7(1).

Ebbes, P., Wedel, M., Böckenholt, U., and Steerneman, T. (2005). Solving and Testing for Regressor-Error (in)Dependence When no Instrumental Variables are Available: With New Evidence for the Effect of Education on Income. *Quantitative Marketing and Economics*, 3(4):365–392.

Ehrlich, I. (1996). Crime, Punishment, and the Market for Offenses. *Journal of Economic Perspectives*, 10(1):43–67.

Elliott, A. (2003). Suit Accuses Macy's of Bias Against Minority Shoppers. *The New York Times*.

Evett, S. R., Hakstian, A.-M. G., Williams, J. D., and Henderson, G. R. (2013). What's Race Got to Do with It? Responses to Consumer Discrimination: Responses to Consumer Discrimination. *Analyses of Social Issues and Public Policy*, 13(1):165–185.

Fang, H. and Moro, A. (2011). Theories of statistical discrimination and affirmative action: A survey. In Benhabib, J., editor, *Handbook of Social Economics. Vol. 1A*. Elsevier, North-Holland.

Farivar, C. (2013). Here's what an actual "six strikes" copyright alert looks like.

Farmer, R. and Glass, B. (2010). *Building Web Reputation Systems*. "O'Reilly Media, Inc.". Google-Books-ID: MXJ5iB7BS_0C.

Fombrun, C. (1995). *Reputation: Realizing Value from the Corporate Image*. {Harvard Business School Press}.

Fombrun, C. and Shanley, M. (1990). What's in a Name? Reputation Building and Corporate Strategy. *ACAD MANAGE J*, 33(2):233–258.

Fryer, R. G. (2007). Belief flipping in a dynamic model of statistical discrimination. *Journal of Public Economics*, 91(5-6):1151–1166.

Fudenberg, D. and Levine, D. K. (1989). Reputation and Equilibrium Selection in Games with a Patient Player. *Econometrica*, 57(4):759.

Fudenberg, D. and Levine, D. K. (1992). Maintaining a Reputation when Strategies are Imperfectly Observed. *The Review of Economic Studies*, 59(3):561.

Gardner, E. (2011). Massive 'Expendables' Piracy Lawsuit Dropped But Will Be Refiled Soon (Exclusive).

Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-Based Approaches to Calculating Marginal Densities. *Journal of the American Statistical Association*, 85(410):398–409.

Gelman, A. and Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statist. Sci.*, 7(4):457–472.

Gill, J. (2007). *Bayesian Methods: A Social and Behavioral Sciences Approach, Second Edition*. CRC Press. Google-Books-ID: Iq_epk4mtM4C.

Gneezy, U. and Rustichini, A. (2000). A Fine Is a Price. *The Journal of Legal Studies*, 29(1):1–17.

Goldenberg, J., Libai, B., and Muller, E. (2001a). Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters*, 12(3):211–223.

Goldenberg, J., Libai, B., and Muller, E. (2001b). Using complex systems analysis to advance marketing theory development: Modeling heterogeneity effects on new product growth through stochastic cellular automata. *Academy of Marketing Science Review*, 2001(9):1.

Goldenberg, J., Libai, B., and Muller, E. (2010). The chilling effects of network externalities. *International Journal of Research in Marketing*, 27(1):4–15.

Gutierrez, L. (2015). Shoppers fighting back against 'shopping while black' profiling with social media, lawsuits. *miamiherald*.

Harris, A.-M. G., Henderson, G. R., and Williams, J. D. (2005). Courting Customers: Assessing Consumer Racial Profiling and Other Marketplace Discrimination. *Journal of Public Policy & Marketing*, 24(1):163–171.

Hayashi, Y. (2018). Senate Votes to End Auto-Loan Anti-Discrimination Rule. *Wall Street Journal*.

Hill, R. P. and Stephens, D. L. (2003). The Compassionate Organization in the 21st Century. *Organizational Dynamics*, 32(4):331–341.

Huff, D. L. (1964). Defining and estimating a trading area. *The Journal of Marketing*, pages 34–38.

Hurvich, C. M., Simonoff, J. S., and Tsai, C.-L. (1998). Smoothing parameter selection in non-parametric regression using an improved akaike information criterion. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(2):271–293.

Jin, G. Z. and Leslie, P. (2009). Reputational Incentives for Restaurant Hygiene. *American Economic Journal: Microeconomics*, 1(1):237–267.

Klewes, P. D. J. and Wreschniok, R. (2009). Reputation capital Building and maintaining trust in the 21st century. In Klewes, J. and Wreschniok, R., editors, *Reputation Capital*, pages 1–8. Springer Berlin Heidelberg.

Koren, J. R. (2016). Toyota will compensate black and Asian borrowers to settle bias investigation. *Los Angeles Times*.

Kreps, D. M. and Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253–279.

Kugelmass, H. (2016). "sorry, i'm not accepting new patients": An audit study of access to mental health care. *Journal of Health and Social Behavior*, 57(2):168–183.

Lam, D., Lee, A., and Mizerski, R. (2009). The effects of cultural values in word-of-mouth communication. *Journal of international marketing*, 17(3):55–70.

Lewbel, A. (1997). Constructing Instruments for Regressions With Measurement Error When no Additional Data are Available, with An Application to Patents and R&D. *Econometrica*, 65(5):1201.

Libai, B., Muller, E., and Peres, R. (2013). Decomposing the value of word-of-mouth seeding programs: Acceleration versus expansion. *Journal of Marketing Research*, 50(2):161–176.

Livingston, J. A. (2005). How Valuable Is a Good Reputation? A Sample Selection Model of Internet Auctions. *The Review of Economics and Statistics*, 87(3):453–465.

Luca, M. and Zervas, G. (2016). Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud. *Management Science*, 62(12):3412–3427.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46.

Lucking-Reiley, D., Bryan, D., Prasad, N., and Reeves, D. (2007). Pennies from ebay: The determinants of price in online auctions. *The journal of industrial economics*, 55(2):223–233.

Lundberg, S. J. and Startz, R. (1983). Private discrimination and social intervention in competitive labor market. *The American Economic Review*, 73(3):340–347.

Maddux, J. E. and Rogers, R. W. (1983). Protection motivation and self-efficacy: A revised theory of fear appeals and attitude change. *Journal of Experimental Social Psychology*, 19(5):469–479.

March, J. and Simon, H. (1958). *Organisations*. Wiley.

Mayzlin, D., Dover, Y., and Chevalier, J. (2014). Promotional Reviews: An Empirical Investigation of Online Review Manipulation. *The American Economic Review*, 104(8):2421–2455.

Mazumdar, A. and Shen, T. (2016). Porn Infringement Battles Keep the Action Going for Both Sides.

McCormick, R. E. and Tinsley, M. (1987). Athletics versus Academics? Evidence from SAT Scores. *Journal of Political Economy*, 95(5):1103–1116.

McDevitt, R. C. (2011). Names and Reputations: An Empirical Analysis. *American Economic Journal: Microeconomics*, 3(3):193–209.

McLaughlin, E. C. (2018). Cafe shut down after protesters enter, chanting, 'Starbucks coffee is anti-black!'.

Mendoza, J. P., Wielhouwer, J. L., and Kirchler, E. (2017). The backfiring effect of auditing on tax compliance. *Journal of Economic Psychology*, 62:284–294.

Merle, R. (2018). Senate poised to roll back watchdog effort to prevent discrimination in auto lending market. *Washington Post*.

Merton, R. K. (1968). The Matthew Effect in Science. *Science*, 159(3810):56–63.

Merton, R. K. (1988). The Matthew Effect in Science, II: Cumulative Advantage and the Symbolism of Intellectual Property. *Isis*, 79(4):606–623.

Mick, D. (2006). Meaning and Mattering Through Transformative Consumer Research. *Advances in Consumer Research*, 33(1):1–4.

Milgrom, P. and Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2):280–312.

Murphy, R. G. and Trandel, G. A. (1994). The relation between a university's football record and the size of its applicant pool. *Economics of Education Review*, 13(3):265–270.

Muth, J. F. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica*, 29(3):315–335.

Narasimhan, C. (1984). A price discrimination theory of coupons. *Marketing Science*, 3(2):128–147.

Novick, M. R. (1965). The Axioms and Principal Results of Classical Test Theory. *ETS Research Bulletin Series*, 1965(1):i–31.

Obloj, T. and Obloj, K. (2006). Diminishing Returns From Reputation: Do Followers Have a Competitive Advantage? *Corporate Reputation Review*, 9(4):213–224.

Ouellet, J.-F. (2007). Consumer racism and its effects on domestic cross-ethnic product purchase: An empirical test in the United States, Canada, and France. *Journal of Marketing*, 71(1):113–128.

Pager, D. and Shepherd, H. (2008). The Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit, and Consumer Markets. *Annual Review of Sociology*, 34(1):181–209.

Park, S. and Gupta, S. (2012). Handling Endogenous Regressors by Joint Estimation Using Copulas. *Marketing Science*, 31(4):567–586.

Peres, R. and Van den Bulte, C. (2014). When to take or forgo new product exclusivity: Balancing protection from competition against word-of-mouth spillover. *Journal of Marketing*, 78(2):83–100.

Pettigrew, S. (2001). Why a Journal of Research for Consumers. *Journal of Research for Consumers*, 33(1):1–4.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review*, 62(4):659–661.

Phippen, J. W. (2017). Philadelphia's Lawsuit Against Wells Fargo. *The Atlantic*.

Podoshen, J. S. (2006). Word of mouth, brand loyalty, acculturation and the American Jewish consumer. *Journal of Consumer Marketing*, 23(5):266–282.

Quillian, L. (2006). New Approaches to Understanding Racial Prejudice and Discrimination. *Annual Review of Sociology*, 32(1):299–328.

Rand, W. and Rust, R. T. (2011). Agent-based modeling in marketing: Guidelines for rigor. *International Journal of Research in Marketing*, 28(3):181–193.

Resnick, P. and Zeckhauser, R. (2002). Trust among strangers in internet transactions: Empirical analysis of eBay' s reputation system. In *The Economics of the Internet and E-commerce*, volume 11 of *Advances in Applied Microeconomics*, pages 127–157. Emerald Group Publishing Limited.

Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Exp Econ*, 9(2):79–101.

RIAA (2013). RIAA - Why We Do What We Do - August 06, 2013.

Rogers, R. W. (1975). A Protection Motivation Theory of Fear Appeals and Attitude Change1. *The Journal of Psychology*, 91(1):93–114.

Rust, R. T. and Bornman, E. O. (1982). Distribution-free methods of approximating nonlinear marketing relationships. *Journal of Marketing Research*, 19(3):372–374.

Rust, R. T. and Donthu, N. (1995). Capturing geographically localized misspecification error in retail store choice models. *Journal of Marketing Research*, 32(1):103–110.

Sag, M. (2015). Copyright Trolling, An Empirical Study. *Iowa Law Review*, (1105-1146 (2015)).

Savage, D. G. (2017). A baker refused to make a cake for a gay couple due to religious beliefs. Supreme Court will rule on the case in fall. *Los Angeles Times*.

Schelling, T. C. (1969). Models of segregation. *The American Economic Review*, 59(2):488–493.

Schroeder, P. (2017). American Express to pay $96 million over discriminatory card terms. *Reuters*.

Schwartz, B. (1989). *Psychology of learning and behavior*. WW Norton & Co.

Silver-Greenberg, J. (2012). American Express Agrees to Refund $85 Million. *The New York Times*.

Simon, H. (1982). ADPULS: An Advertising Model with Wearout and Pulsation. *Journal of Marketing Research*, 19(3):352–363.

Smith, E. B. and Rand, W. (2017). Simulating Macro-Level Effects from Micro-Level Observations. *Management Science*.

Stigler, G. J. (1970). The optimum enforcement of laws. *Journal of Political Economy*, 78(3):526–536.

Tadelis, S. (2002). The Market for Reputations as an Incentive Mechanism. *Journal of Political Economy*, 110(4):854–882.

Tadelis, S. (2003). Firm Reputation with Hidden Information. *Economic Theory*, 21(2/3):635–651.

Tanner, J. F., Hunt, J. B., and Eppright, D. R. (1991). The Protection Motivation Model: A Normative Model of Fear Appeals. *Journal of Marketing*, 55(3):36.

Taylor, M. F. (2012a). Addicted to the Risk, Recognition and Respect that the Graffiti Lifestyle Provides: Towards an Understanding of the Reasons for Graffiti Engagement. *International Journal of Mental Health and Addiction*, 10(1):54–68.

Taylor, P. (2012b). *Hackers : Crime and the Digital Sublime*. Routledge.

Trusov, M., Bucklin, R. E., and Pauwels, K. (2009). Effects of word-of-mouth versus traditional marketing: findings from an internet social networking site. *Journal of Marketing*, 73(5):90–102.

U.S. Copyright Office (2018). Chapter 5 - Circular 92 | U.S. Copyright Office.

Uslu, A., Durmuş, B., and Taşdemir, S. (2013). Word of Mouth, Brand Loyalty, Acculturation and the Turkish Ethnic Minority Group in Germany. *Procedia - Social and Behavioral Sciences*, 99:455–464.

Van der Sar, Ernesto (2016). Top 10 Most Popular Torrent Sites of 2016.

Varian, H. R. (1989). Price discrimination. *Handbook of industrial organization*, 1:597–654.

Weber, M. (2009). *The Theory Of Social And Economic Organization*. Simon and Schuster.

Wilensky, U. (1999). *NetLogo*, volume 4952. Center for Connected Learning and Computer-based Modeling.

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics*, 20(1):39–54.

Wooldridge, J. M. (2015). Control Function Methods in Applied Econometrics. *Journal of Human Resources*, 50(2):420–445.

Yoganarasimhan, H. (2013). The Value of Reputation in an Online Freelance Marketplace. *Marketing Science*, 32(6):860–891.

Zhao, M. and Xie, J. (2011). Effects of social and temporal distance on consumers' responses to peer recommendations. *Journal of Marketing Research*, 48(3):486–496.