informs
http://pubsonline.informs.org/journal/isre

INFORMATION SYSTEMS RESEARCH
Articles in Advance, pp. 1–23
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Evolution of Referrals over Customers' Life Cycle: Evidence from a Ride-Sharing Platform

Carlos Fernández-Loría,[a,*] Maxime C. Cohen,[b] Anindya Ghose[c]

[a] Department of Information Systems, Business Analytics, and Operations Management, HKUST Business School, Hong Kong, China;
[b] Department of Operations Management, Desautels Faculty of Management, McGill University, Montreal, Quebec H3A 1G5, Canada;
[c] Department of Technology, Operations, and Statistics, Stern School of Business, New York, New York 10012
*Corresponding author

Contact: imcarlos@ust.hk, https://orcid.org/0000-0003-4509-3768 (CF-L); maxccohen@gmail.com, https://orcid.org/0000-0002-2474-3875 (MCC); aghose@stern.nyu.edu, https://orcid.org/0000-0002-6499-8944 (AG)

Received: April 7, 2020
Revised: February 16, 2021; November 1, 2021; March 8, 2022
Accepted: April 10, 2022
Published Online in Articles in Advance: July 15, 2022

https://doi.org/10.1287/isre.2022.1138

Abstract. Online platforms often ask their users to refer friends in exchange for a reward. This paper addresses how referral generation and referral value evolve throughout the customer's life cycle as a function of service usage, experience level, and past referral behavior. Our analysis is based on a longitudinal data set that comprises the transactions and referral actions of 400,000 users in a ride-sharing platform over a year. The richness of our data set allows us to address two shortcomings from previous studies: modeling dynamic behavior (i.e., the relationship between past and future referrals) and accounting for unobserved heterogeneity across users. Our results show that users make more referrals when they have used the service recently and intensively. For example, users become 9% less likely to make referrals for each week they have not used the service. Furthermore, users make more high-value referrals as they become more experienced with the service. The referral generation and referral value of the top 10% most experienced users are more than 18% higher relative to when they first used the service. Finally, as users make more referrals, they become more likely to run out of friends to whom they can refer the service, leading to less referrals in the future. After users make their first referral, the probability of making additional referrals decreases by more than 78% and the value of subsequent referrals reduces by 19% on average. The results imply that platforms should consider tailoring their referral programs according to how referral generation and referral value evolve over time.

History: Xiaoquan (Michael) Zhang, Senior Editor; Khim Yong Goh, Associate Editor.

Keywords: referrals • ride-sharing • online platforms

1. Introduction

Viral marketing and the business of influencers have proliferated in recent years. As Mark Zuckerberg, Facebook's CEO, puts it[1]: "People influence people. Nothing influences people more than a recommendation from a trusted friend. A trusted referral influences people more than the best broadcast message. A trusted referral is the Holy Grail of advertising." Empirical evidence supports this claim. According to Nielsen (2012), customers are four times more likely to buy a product from a particular brand when referred by a friend, and 84% of people follow recommendations from family and friends.[2] More recently, referrals have received increasing attention because of the critical role they play in the early development and growth of startups, especially for technology firms (Koch and Benlian 2015) and sharing economy platforms (Dillahunt and Malone 2015).

A large number of studies in the referral literature have been devoted to the empirical assessment of a variety of factors that may induce customers to make referrals. Examples include monetary rewards (Ryu and Feick 2007), the strength of the social tie between the referrer and the referred user (Wirtz et al. 2013), and the referrer's degree of experience with the firm (Van den Bulte et al. 2018). Within this research stream, a series of recent studies have focused on the interplay between referrals and customer behavior. For example, Schmitt et al. (2011a) and Van den Bulte et al. (2018) investigate the effect of customer tenure on referral profitability, whereas Garnefeld et al. (2013) consider the effect of referrals on customer loyalty.

Studying the interplay between referrals and customer behavior is a challenging problem for several reasons. First, randomized experiments in this context are usually impossible (e.g., one cannot randomly manipulate the amount of customer tenure), so most empirical studies on referrals are based on observational analyses that are prone to endogeneity issues (Belo and Li 2018). Second, data restrictions often limit researchers to cross-sectional analyses, making it challenging to underpin the exact nature of the relationship between referrals and other customer behaviors. For example, Van den Bulte et al. (2018) find that higher-margin referrals are associated with referrers who have more experience with the firm. However, the nature of their study does not allow to conclude whether margins become higher as referrers become more experienced. For instance, one alternative explanation to referrers becoming more profitable as they become more experienced is that customers tend to refer the service to individuals with similar characteristics to them, and because experienced customers are more likely to have characteristics that make them a good match to the firm, this would result in experienced customers making more profitable referrals.

A third challenge is that customer behavior is rarely constant over time. Yet, no study has focused on how referral behavior evolves throughout the customer's life cycle. For example, online platforms often offer referral programs asking their users to invite several friends in exchange of free access to premium features (Belo and Li 2018), but no study has analyzed how past referral behavior affects subsequent referrals. Also in the context of an online platform, Jung et al. (2020) show that calls to action encouraging referrals are more effective when targeted to users who make frequent purchases. However, since their analysis is cross-sectional, their results do not necessarily imply that calls to action are more effective *when* users are making frequent purchases. Similarly, Garnefeld et al. (2013) find that newer customers are more likely to participate in referral programs. Nevertheless, one cannot conclude from their study whether customers become less likely to make referrals as time elapses, or if the gap in referral participation should be attributed to unobserved differences between newer and more experienced customers.

We address these limitations from the referral literature by exploiting a rich and comprehensive longitudinal data set on customer behavior. We analyze referrals in the context of an online ride-sharing platform and investigate the underexplored question of how referral generation and referral value evolve over time throughout the customer's life cycle, using data from all the trips and referrals made by 400,000 users over a year. Importantly, observing each user multiple times over an extended time period allows us to model dynamic behavior (such as the relationship between past and future referrals) and account for unobserved heterogeneity across users.

Our results show that users make more referrals when they have used the service recently and intensively, hence supporting prior research recommending the timing of calls to action based on purchasing behavior (Jung et al. 2020). More specifically, we find that users become 13% more likely to make referrals for each ride they complete in a given week and 9% less likely to make referrals for each week they have not used the service. Likewise, referral value increases by 1.6% for each ride a user completes in a given week. Contrary to previous cross-sectional analyses (Garnefeld et al. 2013), we find that users make more and higher-value referrals as they become more experienced with the service: referral likelihood and referral value increase by 2.7% and 2.8%, respectively, for every 10 rides completed by the user. Interestingly, we also retrieve some results from prior studies when not accounting for customer dynamic behavior and unobserved heterogeneity. As part of our analysis, we provide several explanations for the difference in the findings, including potential issues due to time-invariant unobserved heterogeneity and other context-specific characteristics. Finally, we show that as users make more referrals, they become more likely to run out of friends to whom they can refer the service, hence leading to less (and lower value) referrals in the future. We find that, when users make their first referral, the probability of making additional referrals decreases by more than 78%, and the value of subsequent referrals decreases by 19% on average. Thus, our findings suggest that firms should strategically invest more in earlier referrals in the customer journey since they are the most valuable.

Overall, this paper advances our understanding on how referral generation and referral value evolve with time-varying factors that describe customer behavior (service usage, experience level with the service, and past referral behavior). Ultimately, our results offer managerial insights that can be used to tailor referral programs throughout the customer's life cycle.

2. Theoretical Background

Recent studies in the word of mouth (WOM) literature have focused on the design of online referral programs. Key design choices include the incentive design (Hong et al. 2017, Jung et al. 2021, Belo and Li 2018), the call to action for online referrals (Jung et al. 2020), and the design of the message from the person initiating the referral to the person receiving the referral (Sun et al. 2021). Although these studies emphasize practical advice for the general design of referral programs (e.g., "use split rewards" and "use a prosocial framing for the call to action"), a recurrent finding is that the effectiveness of different designs is often moderated by individual-level characteristics. For example, split rewards work better when the social distance between the sender and the recipient is large because the extrinsic motivations behind rewards are more salient for weaker social ties (Hong et al. 2017). A prosocial framing of the call to action works particularly well for customers with a higher affinity with the product (Jung et al. 2020) or for individuals with a large social value orientation (Huang et al. 2019). In a dating platform, Belo and Li (2018) find that asking older, well-educated senders to refer more friends before receiving free access to premium features was effective in increasing the platform growth, while asking the same to female senders reduced their engagement. Overall, these findings suggest that a better understanding of which factors influence referral behavior is a key step toward customized referral programs.

Another important research stream in the referral literature has focused on the factors that may motivate individuals to make referrals. Examples include studying the interplay between referrals and tie strength (Ryu and Feick 2007) and investigating the best timing to offer the reward (Biyalogorsky et al. 2001). Our study builds on this line of work by enhancing our understanding on how the value of a referrer, in terms of his or her referral outcomes, evolves over time and throughout the customer's life cycle. More broadly, the value of a referrer can be decomposed in terms of three critical referral outcomes:

$$\text{Referrer Value} = \text{Referral Attempt Rate} \times \text{Conversion Rate} \times \text{Referral Value}, \tag{1}$$

where the referral attempt rate is the number of times the referrer attempts to refer someone, the conversion rate is the probability that a referral attempt turns into an actual referral, and the referral value is the referred individual's value to the firm. Alternatively, in settings where the referral attempt rate and the conversion rate are not observed directly, the referrer value may be decomposed in terms of two referral outcomes:

$$\text{Referrer Value} = \text{Referral Generation Rate} \times \text{Referral Value}, \tag{2}$$

and the referral generation rate is the referral attempt rate multiplied by the conversion rate.

We study how the referrer value evolves over time by focusing on the interplay between the above key referral outcomes and several time-varying factors, such as service usage, experience level with the service, and past referral behavior. These factors evolve over time in the customer's life cycle and can be accurately measured by online platforms to subsequently design better referral programs and decide on the best timing to send out calls for referrals.

The study of time dynamics between customer behavior and desirable outcomes has received substantial attention in the marketing literature. For example, prior studies have looked at the interplay (over time) between service usage and customer satisfaction (Bolton and Lemon 1999), customer retention (Ascarza and Hardie 2013), and future purchases (Lemon and Wangenheim 2009). These studies have shown that time dynamics can have important implications for the design of pricing strategies, loyalty programs, and customer retention programs. For example, time dynamics can be used to inform cross-selling promotions (Lemon and Wangenheim 2009) and to identify the stages customers go through before cancellation to optimize marketing interventions (Ascarza and Hardie 2013). More generally, a better understanding of how customer behavior evolves over time can help companies design better marketing interventions and make improved targeting decisions (Chen et al. 2005, Zhang et al. 2015).

A few studies in the referral literature have also looked at longitudinal variations in customer behavior. For example, Schmitt et al. (2011a) and Van den Bulte et al. (2018) examine whether the benefits of referrals erode over time. Specifically, Schmitt et al. (2011a) assess the relationship between referrals and customer value (in terms of profitability and retention rate), whereas Van den Bulte et al. (2018) explore the underlying mechanisms of social enrichment and better matching to explain why referred customers exhibit higher margins and lower churn rates relative to customers acquired via traditional channels. Additionally, Garnefeld et al. (2013) find that participation in a referral program can increase customer loyalty. Nevertheless, these studies focus primarily on how referrals may affect the profitability of customers rather than on how referral outcomes (as defined in Equations (1) and (2)) evolve over time.

A unique strength of our study relative to prior empirical analyses in the referral literature—most of which use cross-sectional data or a single field experiment—is our comprehensive data set on longitudinal variations of customer referrals. More specifically, we observe transactions and referrals on a weekly basis during one year for 400,000 users. Thus, because our data include several occurrences of the same customers making referrals, our models can account for unobserved (time-invariant) heterogeneity in customers, allowing us to analyze the relationship between referrals and time-varying factors. Moreover, our data allow us to incorporate dynamics in customer behavior, such as the interplay between past and future referrals, which has not been considered in previous studies. Therefore, the granularity and scale of our data allow us to (i) address questions related to the customer's life cycle that cannot be answered using cross-sectional data or a single field experiment, and (ii) control for unobserved heterogeneity across users to strengthen or refute the findings from previous studies. We formally discuss our contributions relative to prior studies in Section 2.1 and summarize them in Table 1.

2.1. Hypothesis Development

In this section, we develop hypotheses on how referral outcomes evolve over time and throughout the customer's life cycle. More specifically, we focus on assessing

Table 1. Summary of Empirical Findings and Research Contributions

	Temporal relationship with		
	Referral generation	Referral value	Research contributions
Current service usage	+	+	First temporal (rather than time-fixed) empirical assessment. Attribution to time-varying mechanisms rather than homophily.
Experience level	+	+	Resolution of extant incongruencies in the referral literature. First temporal (rather than time-fixed) empirical assessment. Attribution to time-varying mechanisms rather than homophily.
Past referrals made	—	+	First assessment of the relationship between past and future referrals.
Past referral value	—	—	

how referral generation and referral value (Equation (2)) evolve with (i) users' current service usage, (ii) users' experience level with the service, and (iii) users' past referral behavior. We formally define all these quantities in Section 3 and propose here several latent mechanisms that can explain our empirical results. Such mechanisms include customer satisfaction, customer loyalty, knowledge about the firm's offerings and referral recipients, social enrichment, and the customer's pool of potential referrals.

Figure 1 provides an illustration of the framework underlying our hypotheses. The boxes with solid lines are the variables of interest, whereas the boxes with dotted lines represent the latent mechanisms through which these variables relate. The dashed arrows represent positive associations that have been used in the literature for the operationalization of the latent variables under consideration. The solid arrows represent causal relationships.

2.1.1. Referral Generation. The fundamental building block of successful WOM-based product diffusion is delight or satisfaction among existing customers (Kornish and Li 2010), which is then communicated to relevant parts of their social networks that may also experience a similar satisfaction from adopting the focal product (Anderson 1998). Thus, given that service usage is both an antecedent and a consequence of customer satisfaction (Bolton and Lemon 1999), an increase in service usage should have a positive association with referrals.

Hypothesis 1a. *Users become more likely to refer when they are using the service intensively.*

It is known that customer satisfaction and the enthusiasm to share the product may decay with the amount of time following the last purchase (Berger and Schwartz 2011). Furthermore, a long inactive period can be indicative of the fact that the customer may have permanently left the platform. We thus hypothesize the following complementary hypothesis.

Hypothesis 1b. *Users become less likely to refer when they have not used the service recently.*

The extant referral literature is ambivalent regarding the relationship between experience level and referral behavior. Garnefeld et al. (2013) find a negative correlation between customer tenure, which is used as a measure of experience level in their study, and participation in referral programs. The authors argue that one possible explanation for this negative correlation is that customers are more likely to articulate recommendations right after they have adopted the service because newer customers are more likely to communicate the "goodness" of their choice to others, either to convince themselves or to prevent others from disregarding their ability to make good choices (Wangenheim and Bayón 2007).

In contrast, there are two alternative theoretically motivated mechanisms that may result in a positive association between experience level and referral behavior. The first is related to customer loyalty, which is often associated with various measures of the customer's experience level, such as customer tenure and the number of repeated purchases (Yi 1990, Hallowell 1996). Importantly, loyal customers with a longer relationship with the firm may sense more satisfaction and benevolence toward the firm, and as a result exert a greater effort in making referrals (Yi 1990, Van den Bulte et al. 2018). The second mechanism is active matching, which involves deliberate screening and occurs when current customers know the firm's offerings better than noncustomers and selectively match some of their peers to the firm (Van den Bulte et al. 2018). Thus, because experience leads customers to have a better understanding of the firm's offerings, the experience level should positively influence referrals via active matching. Together, these two mechanisms imply the following.

Hypothesis 1c. *Users become more likely to refer when they are more experienced with the service.*

Although we are not aware of any study empirically analyzing the relationship between past and future referrals, it is clear that the referral potential of individuals is determined by the size of their social networks (Helm 2003). Thus, we should expect that as customers make more referrals, they also become

Figure 1. Research Framework and Hypotheses

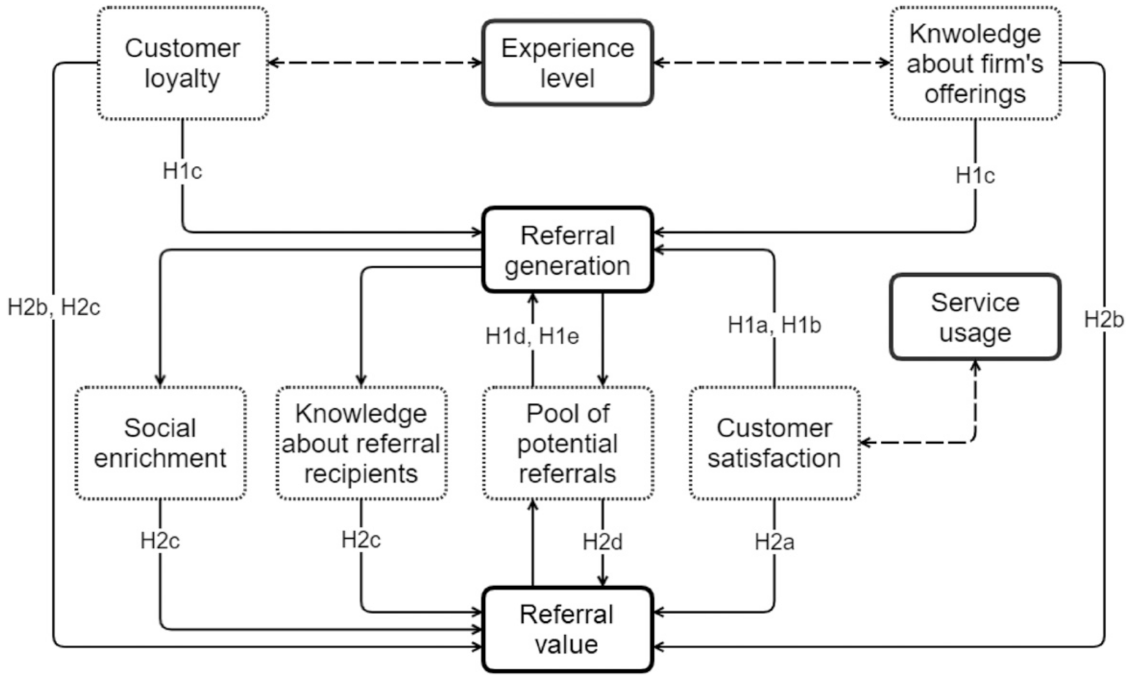

more likely to "run out" of friends to whom they can refer the service.

Hypothesis 1d. *Users become less likely to refer when they have made more referrals in the past.*

Given that the probability of purchasing as a result of a referral is higher when the sender and the recipient are similar (Wangenheim and Bayon (2007)), higher-value referrals are likely to correspond to the strongest social ties of the referrer. However, the number of such "good matches" (strong ties that appreciate the service) in an individual's social network is limited (Helm 2003)). When users refer the service to their stronger ties, the remaining individuals in their network correspond to weaker ties that are less likely to accept the referral. Therefore, we hypothesize the following.

Hypothesis 1e. *Users become less likely to refer when the value of their past referrals is higher.*

2.1.1.1. Summary. Users are more likely to make referrals when they have used the service recently and intensively and when they become more experienced with the service. We propose the following underlying mechanisms for these hypotheses: customer satisfaction, customer loyalty, and knowledge about the firm's offering. In contrast, customers become less likely to refer as they make referrals, particularly if the past referrals were of high value. We argue that this is the result of a reduced pool of potential people to refer and because higher-value referrals correspond to stronger ties in the referrer's social network, so the remaining, weaker ties are less likely to accept referrals.

2.1.2. Referral Value. Helm (2003) discusses customer satisfaction as a moderator of referral value. The reasoning is that greater satisfaction is what leads customers to advocate for (as opposed to just talk about) products (Buttle 1998). Therefore, users who make referrals because they are highly satisfied with the service will do so differently than opportunistic users who just care about referral rewards; the former are more likely to genuinely care about convincing others to use the service. In our research setting, higher usage implies more value to the firm. Thus, we hypothesize the following.

Hypothesis 2a. *Referral value increases when users are using the service intensively.*

In the context of financial services, Van den Bulte et al. (2018) find that more experienced referrers generate higher-margin referrals (i.e., higher-value referrals). The authors suggest three theories to explain this finding: one theory that involves a time-invariant correlational association and two theories that rely on time-varying mechanisms. The first (time-invariant) theory is that experienced referrers have unobservable characteristics that make them a good match to the firm, and their referred friends often share these characteristics. This is also known as passive or homophily-

Published in Information Systems Research on July 15, 2022 as DOI: 10.1287/isre.2022.1138.
This article has not been copyedited or formatted. The final version may differ from this version.

based matching. The second theory relies on active matching: since experienced customers tend to have a better understanding of the firm's offerings, they will generate better matches via referrals. The third theory relates to customer loyalty: to the extent that customers with a longer relationship with the firm also feel more satisfaction toward the firm, they will be less likely to generate referrals opportunistically just to earn the reward (Jing and Xie 2011, Schmitt et al. 2011b). Therefore, assuming that the positive association between experience level and referral value is driven by the time-varying associations described by the latter two theories, we hypothesize the following.

Hypothesis 2b. *Referral value increases when users become more experienced with the service.*

There are theoretical reasons to believe that the act of referring someone could also improve (or be positively associated with) the value of future referrals. First, the marketing literature has long considered product recommendations as a signal of customer loyalty (Yi 1990, Hallowell 1996). Thus, to the extent that the number of previous referrals serves as a measure of customer loyalty, referral value should increase with more referrals.

Another possibility is that referral value improves through social enrichment, which relates to the fact that the relationship between the referred customer and the firm is enriched by a common third party. Social enrichment theory was primarily developed in the context of labor economics and employee referrals (Fernandez et al. 2000, Castilla 2005) but has also been used to support empirical findings in referral studies (Van den Bulte et al. 2018). In essence, this theory asserts that the bond between a customer and the firm is strengthened by the presence of a known third party. Since it is common for users to make referrals to family members and coworkers, the presence of a fellow customer may provide social or functional benefits, such as education and discussion about the service. Consequently, if a newly referred customer also knows some of the previously referred customers, social enrichment implies that the number of previous referrals could positively affect the value of future referrals.

Finally, another possible mechanism, related to active matching, is that the referrer may learn to make better matches as a consequence of previous referrals, that is, learning by doing (Galenianos 2013). For example, a referrer may learn that a friend to whom he or she referred the service really liked it, and subsequently refers the service to similar friends. Therefore, we hypothesize the following.

Hypothesis 2c. *Referral value increases when users have made more referrals in the past.*

On the other hand, the value of future referrals will be lower when users have already referred the service to the best matches in their social network. It is well known that users are more likely to refer the service first to strong social ties (Brown and Reingen 1987, Krackhardt et al. 2003, Sun et al. 2019), particularly those that are more likely to appreciate the service. As a result, the potential value of the remaining individuals in the users' social network is likely to be lower.

Hypothesis 2d. *Referral value decreases when the value of the users' past referrals is higher.*

2.1.2.1. Summary. Referral value increases when customers (i) use the service intensively, (ii) become more experienced with the service, and (iii) make more referrals. However, (iv) referral value is lower when the customer's past referrals are of high value. We propose the following underlying mechanisms for these hypotheses: customer satisfaction for (i), customer loyalty and active matching—that is, knowledge about the firm's offering and potential recipients—for (ii) and (iii), social enrichment for (iii), and finally, changes to the pool of potential referrals for (iv).

Hypothesis 2c and Hypothesis 2d imply that past referral behavior can have opposing effects on referral value: on the one hand, referral value is higher when the customer has made more referrals in the past, but on the other hand, referral value is lower when those past referrals are of high value. Which relationship is stronger is an empirical question that we address in Section 5.2. We find that, in general, the relationship between past referral behavior and future referral value is negative, which means that Hypothesis 2d is more empirically relevant than Hypothesis 2c.

Table 1 presents a summary of the empirical findings and theoretical contributions of our study.

3. Research Setting

We analyze referrals in the context of an online ridesharing platform, referred to as the platform.[3] In the referral program used by the platform, existing users can recommend the service to their friends via one of the following channels: sending a text message or email, sharing a code or a link on social media platforms, or directly communicating a referral code. Then, the friend can use the referral to try the service while earning a reward (the referrer and the referred user each receive $10 worth of credit). The credit can only be spent in rides through the platform, does not expire, and is awarded only when the referred user completes his or her first ride. For many platforms (including our industry partner), a significant portion of accounts originate from referrals.

As mentioned, this study is about understanding how referral outcomes evolve over time and throughout the customer's life cycle. We focus specifically on two types of referral outcomes: (1) the probability that a user makes a successful referral (referral generation) and (2) the value of the referral. We formally define these two types of referral outcomes next.

A successful referral occurs when the referred customer enters the referral code in the smartphone application (either manually or by using a link received via email or text message). Consequently, our main analysis focuses on referrals that materialized and does not incorporate referral requests that were sent but not accepted by the referred user. This limitation is very common (both in practice and in the academic literature) due to the challenge of accurately tracking referral requests that were not accepted (Sun et al. 2021). We highlight that we still include the observations from customers who accepted a referral (i.e., created an account and entered the referral code) but did not complete a ride. In Section 6.4, we conduct additional analyses to address this limitation by using a different data set that captures a large portion of referral attempts, allowing us to further investigate the relationship between referral attempts and referral generation.

We define referral value in terms of the number of rides completed by the referred customer after account creation. As discussed in more detail in the following sections, we consider the number of rides completed in a time window of 13 weeks (i.e., three months) in our analysis, but our qualitative findings remain the same when using other time windows (e.g., 8 and 16 weeks).

3.1. Data

Our data consist of a random sample of 400,000 user accounts created in 2017 in three cities where the ride-sharing platform is operating. The data include all the rides and referrals that each of these users made in a period of 51 weeks (from January 2 to December 24). We excluded all users who made more than 10 referrals (there were 377 such accounts, less than 1% of all referrers). Our industry partner believes that such accounts are not representative and can be connected to unsanctioned marketing activities, such as posting a referral code on social forums. Indeed, these 377 referrers made more than 40 referrals on average (many of them did several hundred referrals).

We use the data described above to create two panel data sets. The first panel data set is defined at the week-user level and is used to model the probability that a user makes a successful referral in a given week (i.e., referral generation). This panel comprises 9,527,831 week-user rows. Each row contains the following fields: user ID, week number, number of rides completed by the user during that week (*rides*), number of accumulated past rides at the beginning of the week (*rides_before*), city (*city*), number of weeks since the user took a ride or created his or her account (*inactive_t*), number of past referrals (*prev_refs*), and average number of rides completed by the user's referral recipients in their first 13 weeks (*prev_value*),[4] Moreover, each row includes whether the user made a successful referral in the next week (*ref*), which is the dependent variable.[5]

The second panel data set is defined at the referral-referrer level and is used to model referral value. This data set includes 69,597 referrals made by 42,195 individuals. Each row in this data set contains the same fields discussed above at the time the referral was made: user ID, week number, number of rides completed in the last seven days (*rides*), number of accumulated past rides (*rides_before*), city (*city*), number of weeks since the referrer completed a ride or created his or her account (*inactive_t*), number of past referrals (*prev_refs*), and value of those past referrals (*prev_value*). Additionally, each row includes the number of rides completed by the referred customer in the next 13 weeks (*value*), which is the dependent variable in this case.

4. Empirical Strategy

In this section, we present our empirical models. In Section 4.1, we consider referral generation, whereas in Section 4.2 we consider referral value.

4.1. Referral Generation: Model

We model the probability of making a successful referral in a given week using a logistic regression model—a standard approach when regressing binary outcomes (Harrell 2015) that has also been previously used in the context of referrals and customer acquisition (see Wangenheim and Bayón 2007 for a concrete example and an overview of related studies). To account for time-invariant unobserved heterogeneity in customers (e.g., demographics, income level, neighborhood), we include fixed effects at the user level. Although incorporating fixed effects in a logistic regression model forces us to drop observations from users who never made a referral, we still retain 918,749 observations. Furthermore, the qualitative results remain the same when the entire sample is used to estimate a linear probability model with fixed effects (for more details, see Section 6.1). Our model is the result of estimating the following logistic regression equation:

$$P(ref_{i,t+1} = 1) = F(\alpha_i + \gamma_t + \boldsymbol{\beta_1 X_{it}} + \beta_2 prev\_refs_{it} + \beta_3 prev\_value_{it}), \quad (3)$$

where F is the logistic cumulative distribution function and the dependent variable, $ref_{i,t+1}$, equals one if customer i made a referral the week after week t and zero otherwise. This formulation ensures that all the variables used to estimate future referrals are measured before the referrals are made. All the independent variables used in the model are defined in Section 3.1.

The vector of variables $\boldsymbol{X_{it}}$ in Equation (3) consists of $rides\_before_{it}$, $inactive\_t_{it}$, and $rides_{it}$, which are similar to the RFM (recency, frequency, and monetary) metrics commonly used in the marketing literature (Wei et al. 2010). More precisely, we use these variables to operationalize service usage and experience level as follows:

1. Current usage intensity is measured by $rides_{it}$ (to test Hypothesis 1a).
2. Service usage recency is measured by $inactive\_t_{it}$ (to test Hypothesis 1b).
3. Experience level is measured by $rides\_before_{it}$ (to test Hypothesis 1c).

We operationalize past referral behavior using $prev\_refs_{it}$ (to test Hypothesis 1d) and $prev\_value_{it}$ (to test Hypothesis 1e), which respectively capture past quantity and value. Incorporating dynamics between past and future referrals is particularly important to study the relationship between experience level and the probability of making a referral. Ignoring this dynamics may induce a downward bias because experienced customers are more likely to have made referrals in the past (because they have been using the service for longer), and users who made referrals in the past have less remaining friends to whom they can refer the service. Thus, omitting these variables could produce a spurious association between experience level and referrals. We illustrate this phenomenon in Section 5.1.

Finally, we control for unobserved heterogeneity using fixed effects α_i at the user level, and we include monthly time fixed effects denoted γ_t. We do not use weekly time fixed effects in the main specification of our model because of the high multicollinearity between the week dummies and other time-varying variables. For robustness purposes, we consider the alternative specification with weekly-time fixed effects in Section 6. We find that the main results are robust to the inclusion of time fixed effects either at the weekly, biweekly, or monthly level (see Section 6.1). In Section 6.1, we also consider including alternative time-varying control variables, such as weekly total demand for ride-sharing (across all platforms), and instrumental variables based on weather conditions. These alternative specifications are meant to control for potential sources of time-varying confounding.

Empirical researchers typically deal with dynamics in behavior using models that incorporate lagged variables as predictors. Standard linear models include the Arellano–Bond estimator (Arellano 2003), and for structural equation models see, for example, Allison et al. (2017). However, the nonlinear nature of the logistic model prevents us from directly applying these methods. Instead, we use the maximum-likelihood (ML) estimator applied to a dynamic logistic model with fixed effects proposed by Stammann et al. (2016).

When dealing with binary choice models with fixed effects, it is common to use the conditional maximum-likelihood (CML) estimator instead of the ML estimator because the latter suffers from the incidental parameters problem (Neyman and Scott 1948) when the number of observations per individual is small—leading to biased estimates (see Arellano and Hahn 2007, for a comprehensive review). However, ML is more appropriate in our context for several reasons. First, CML with endogenous lagged variables and fixed effects (Bartolucci and Nigro 2012) is often model specific and imposes the Markov assumption that the dependent variable is conditionally independent of its value in earlier periods given the previous period. In our setting, however, the dependent variable is likely to depend on the total number of referrals made in the past, and not only on whether a referral was made in the previous period. Second, we are unlikely to suffer from the incidental parameters problem given that we have a large number of observations per individual (we observe most of the accounts several dozens of times), and we incorporate the bias correction suggested by Hahn and Newey (2004).[6] Third, CML does not estimate individual fixed effects, which makes it impossible to obtain average partial effects and provide economic interpretations.

4.2. Referral Value: Model

We model referral value using count data models (recall that our dependent variable in this case is the number of rides). With the exception of the Poisson model, count data models with fixed effects are known to suffer from the incidental parameters problem (Lancaster 2000, Greene 2007). Thus, because we want to include fixed effects at the referrer level to account for time-invariant unobserved heterogeneity in customers, we estimate a Poisson regression for the number of rides completed by the referred customer. Previous studies that considered count data models with fixed effects also used a similar approach (Ghose and Han 2011). Robustness tests to address potential issues due to overdispersion are presented in Appendix B.

Incorporating fixed effects at the referrer level implies that the observations from referrers who made a single referral will be dropped; this is an issue for any model that incorporates fixed effects because fitting fixed effects for users who have a single observation is equivalent to dropping these observations. Thus, one limitation of our empirical approach to model referral value is that the findings may not generalize to users who made a single referral. Nevertheless, we are not aware of any observational data set or modeling technique that could address this limitation while accounting for time-invariant confounding, and we still have 39,595 referrals originating from 13,181 referrers after dropping the observations from users who made a single referral.

To model referral value, we estimate the following Poisson regression specification:

$$value_{ij} \sim \text{Poisson}(\lambda_{ij}),$$
$$\lambda_{ij} = \exp(\alpha_i + \gamma_t + \boldsymbol{\beta_1 X_{ij}} + \beta_2 prev\_refs_{ij} + \beta_3 prev\_value_{ij}), \quad (4)$$

where the dependent variable, $value_{ij}$, corresponds to the number of rides completed by customer j (the referred account) during the following 13 weeks after being referred by customer i (the referrer). As discussed, our findings remain consistent when using other time windows (e.g., 8 and 16 weeks). As in Equation (3), the set of variables $\boldsymbol{X_{ij}}$ in Equation (4) consists of $rides\_before_{ij}$, $inactive\_t_{ij}$, and $rides_{ij}$, and we also incorporate the number of previous referrals made by customer i ($prev\_refs_{ij}$) and the average value of previous referrals made by customer i ($prev\_value_{ij}$) as predictors. This specification is meant to test Hypotheses 2a, 2b, 2c, and 2d.

Finally, unobserved heterogeneity at the referrer level is controlled by α_i, and we include monthly time fixed effects denoted γ_t.[7] In line with other studies that used dynamic count data models with fixed effects (Ghose and Han 2011), we estimate our models using the generalized method of moments (GMM). In Section 6.3, we also present an alternative specification that combines referral generation and referral value under the same unit of analysis.

5. Results

In this section, we present the main results of estimating our models for referral generation (Section 5.1) and referral value (Section 5.2).

5.1. Referral Generation: Results

Before discussing our main results for referral generation, we first elaborate on the importance of accounting for unobserved heterogeneity and dynamic behavior when assessing how referral behavior evolves throughout the customer's life cycle. We do so by comparing our fully specified models to simpler versions. Table 2 reports such comparisons for our logistic regression model in Equation (3). Specifically, Column M1 shows the results without fixed effects and without dynamic variables (i.e., *prev_refs* and *prev_value* are not included), Column M2 without dynamic variables, and Column M3 without fixed effects. Finally, Column M4 shows the results of the fully specified model. Results that change from one specification to another are highlighted in bold.

As shown in Table 2, the variables *rides* and *inactive_t* are consistent across all specifications; that is, the coefficient sign does not change and remains statistically significant in all columns. However, this is not the case for *rides_before*: this variable has a negative coefficient in Columns M1–M3 and a positive coefficient in M4. Similarly, the coefficients for *prev_refs* and *prev_value* switch from positive (in M3) to negative (in M4). Thus, the results for these variables are sensitive to the model specification and, hence, fixed effects and dynamic variables make a difference.

Regarding experience level, the results in M1 suggest that experienced users are less likely to make referrals relative to inexperienced users, a finding previously observed by Garnefeld et al. (2013). However, the results from their model do not conclude anything on how the referral behavior evolves as customers become more experienced. Customers with different levels of experience are likely to be different in unobserved ways, and in particular in ways that made

Table 2. Comparison of Logistic Regression Models

	M1 Pooled	M2 FE	M3 Dynamic	M4 FE, dynamic
rides_before	−0.00339***	−0.00824***	−0.00915***	**0.00235*** **
	(−11.57)	(−22.50)	(−25.03)	**(6.92)**
rides	0.120***	0.0100***	0.128***	0.130***
	(32.03)	(36.64)	(37.59)	(43.60)
inactive_t	−0.208***	−0.0768***	−0.193***	−0.104***
	(−85.72)	(−48.01)	(−86.34)	(−67.97)
prev_refs			0.444***	**−1.590*** **
			(82.83)	**(−144.56)**
prev_value			0.0120***	**−0.0454*** **
			(14.47)	**(−28.11)**
N	9,527,831	918,749	9,527,831	918,749
Users	400,000	30,917	400,000	30,917
Pseudo-R^2	0.111	0.107	0.136	0.204
User fixed effects	No	Yes	No	Yes

Notes. Sign changes are highlighted in bold. All models include monthly time fixed effects. Standard errors are adjusted by clustering at the user level.
$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. t Statistics in parentheses.

them adopt the service earlier and stay with the firm. Since such unobserved characteristics could be the main reason behind the negative correlation, their model does not allow us to assess the relationship between experience level and referrals over time.

Incorporating fixed effects partially addresses this issue by controlling for time-invariant unobservable characteristics; but even then, the model may still suffer from confounding bias. For example, M2 includes fixed effects but does not account for the fact that experienced users are more likely to have made past referrals (since they have been using the service for longer). Given that users only know a finite number of people, having made referrals in the past may negatively affect their chances to refer in the future, resulting in *rides_before* having a negative coefficient because of a downward bias. This implies that we need to account for the dynamic behavior to assess the relationship between experience level and referral behavior. Interestingly, only controlling for past referrals (without fixed effects) is not enough, as evidenced by the negative coefficient in M3.

Regarding the relationship between past and future referrals, the positive coefficients for *prev_refs* and *prev_value* in M3 imply that customers who made past referrals are more likely to make future referrals (and even more so when those referrals were of high value). However, this correlation can be attributed to unobserved heterogeneity: past referrers are potentially more likely to know other people interested in the service. Consequently, M3 is useful to predict which users are more likely to make referrals based on their past referral behavior, but not to understand how their referral behavior evolves over time. Critically, one could expect past referrals to have a negative impact on the probability of making a new referral (because the pool of potential candidates decreases). This factor is not captured by the model without fixed effects (M3) and is only present when we incorporate both fixed effects and the dynamic behavior, as in M4, which ultimately allows us to obtain a more accurate picture of the relationship between experience level, past referrals, and future referrals over time.

Our fully specified model (M4) shows a positive coefficient for *rides_before*, which implies that customers make more referrals as they become more experienced with the service (Hypothesis 1c is supported). As mentioned before, this statement is not about comparing recent vs. experienced customers, but rather about how customers' behavior evolves as they become more experienced. Thus, even if more recent customers are more likely to participate in referral programs (either because of unobserved characteristics or because they have a larger pool of friends they can refer), our analysis shows that users are more likely to make referrals as they become more experienced with the service. Similarly, the positive coefficient for *rides* and the negative coefficient for *inactive_t* imply that users make more referrals when using the service intensively and recently (Hypothesis 1a and Hypothesis 1b are supported). Finally, the negative coefficients for *prev_refs* and *prev_value* show that users make less referrals when they have made referrals in the past, particularly if those referrals were of high value (Hypothesis 1d and Hypothesis 1e are supported).

A common issue with logistic regression models is that, due to their nonlinear nature, it is difficult to provide an economic interpretation of the results by just inspecting the coefficients. At the most, one can interpret negative signs as negative correlations and positive signs as positive correlations. In settings such as ours, the issue is further exacerbated because binary choice models with fixed effects treat individual-level effects as nuisance parameters and, hence, are not estimated. However, these parameters are necessary to estimate the magnitude of the effects of interest. Fortunately, the ML estimator we used to estimate our models computes the individual-level parameters, allowing us to compute marginal effects for each variable of interest.

Before discussing these results, we would like to clarify that the marginal effects we report do not necessarily correspond to causal effects. In statistical modeling, partial (or marginal) effects correspond to the partial derivative of the expected value of the outcome variable with respect to a specific regressor. In other words, it is simply the observed increase in the dependent variable when an independent variable increases. For instance, our analysis shows that referral generation increases as referrers become more experienced, but this does not imply that referral generation is increasing *because* users are becoming more experienced. Instead, the theoretical explanation behind our empirical finding is that referral generation increases because users become more loyal or knowledgeable, and these factors are correlated with experience level.

The average marginal effects are reported in Table 3, both in absolute and relative terms. The absolute effect in Table 3 is the average *percentage-point* change in the probability of making a referral when the focal variable increases by one unit. For example, when users make an additional referral (i.e., *prev_refs* increases by one), their probability of making another referral in subsequent weeks decreases by 3.73 percentage points on average. Meanwhile, the relative effect in Table 3 is the average *percentage* change in the probability of making a referral. Therefore, for a user who has a 5% probability of making a referral in a given week, an absolute effect of –3 percentage points implies a relative effect of –60%. Appendix A includes an extended analysis on marginal effects.[8]

Table 3. Average Marginal Effects for Model M4 in Table 2 (Referral Generation)

	Absolute effect	Relative effect
rides_before	0.01 pp***	0.27%
rides	0.56 pp***	13.10%
inactive_t	−0.41 pp***	−9.46%
prev_refs	−3.73 pp***	−78.37%
prev_value	−0.18 pp***	−4.26%

Note. The absolute effect is measured in percentage points.
$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

These results suggest that users are particularly likely to generate referrals when (i) they have used the service intensively and recently, and (ii) they have not exhausted the pool of friends they can refer. For example, users are on average 13% more likely to make referrals for each additional ride they complete in a given week and 9% less likely to make referrals for each additional week they are inactive. The interpretation of our results is more nuanced when *rides* changes from zero to one during a period of inactivity (*inactive_t* > 0). In such circumstances, the marginal effect of an additional ride should also consider the increase in the probability of making a referral that is associated with the termination of the inactivity period when *inactive_t* is set to zero. This implies that the first ride has a higher marginal effect than subsequent rides. Similarly, the marginal effect of the first referral is higher than that of subsequent referrals because a change in *prev_refs* also implies a change in *prev_value* (from zero to the value of the first referral) when users have not made referrals in the past. Thus, the first referral is likely to decrease the probability of making subsequent referrals by more than 78% on average.

Experience level also plays a role in referral generation but is not as economically significant as the two other factors. Nevertheless, one should consider the cumulative impact of experience level when analyzing these estimates. At first, the average relative effect of increasing the experience level by one additional ride may seem small (0.27%), but the implications for users who have been using the platform for a while can be considerable. For instance, around 10% of the users in our sample completed more than 67 rides by the end of 2018. On average, users become approximately 18% more likely ($67 \times 0.27\%$) to make referrals after accumulating such level of experience.

5.2. Referral Value: Results

As with referral generation, accounting for unobserved heterogeneity and dynamic behavior is also important to understand how referral value evolves throughout customers' life cycle. Table 4 compares our fully specified Poisson model in Equation (4) with simpler versions. Column M5 shows the results without fixed effects and without dynamic variables (i.e., *prev_refs* and *prev_value* are not included), Column M6 without dynamic variables, Column M7 without fixed effects, and finally, Column M8 shows the results of the fully specified model. When estimating the models with fixed effects (M7 and M8), we removed the observations from a small group of outliers (21 users, i.e., 0.05% of the users in our sample) who were inactive for more than three months but made high-value referrals during this inactivity period.[9] Results that change signs from one specification to another are highlighted in bold.

In this comparison, the results also vary significantly across specifications: *inactive_t* stops being significant once fixed effects are included (M6 and M8), *rides_before* and *rides* are not statistically significant in M6, and the signs of the coefficients for *prev_refs* and *prev_value* are opposite in M7 and M8. Thus, results

Table 4. Comparison of Poisson Regression Models

	M5 Pooled	M6 FE	M7 Dynamic	M8 FE, dynamic
rides_before	0.00140*	0.000985	0.00171**	0.00282**
	(3.23)	(1.17)	(4.24)	(3.18)
rides	0.0529***	−0.0107	0.0514***	0.0162*
	(11.21)	(−1.37)	(10.93)	(2.13)
inactive_t	−0.0216***	−0.00572	−0.0216***	0.000520
	(−6.81)	(−0.92)	(−6.89)	(0.10)
prev_refs			−0.0633***	**0.038***
			(−7.16)	**(2.44)**
prev_value			0.00708***	**−0.0534*****
			(5.30)	**(−12.24)**
N	69,597	39,516	69,597	39,516
Groups	42,195	13,160	42,195	13,160
Pseudo-R^2	0.026	0.440	0.028	0.515
User fixed effects	No	Yes	No	Yes

Notes. Sign changes are highlighted in bold. All models include monthly time fixed effects. Standard errors are adjusted by clustering at the referrer level.
$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. t Statistics in parentheses.

are once again sensitive to the model specification, so that fixed effects and dynamic variables make a difference.

The model in M7 (i.e., when we do not control for unobserved heterogeneity) shows that users with fewer past referrals make higher-value future referrals (and even more so when past referrals were of high value). Nonetheless, these results may be driven by confounding factors rather than dynamic behavior. First, referrals from users who made many past referrals could be of lower value because such users are potentially sending a large number of referral requests rather than carefully targeting individuals who are likely to appreciate the service. Second, users who made higher-value referrals in the past may be more likely to know other potentially good customers. Thus, controlling for unobserved heterogeneity is critical in order to understand the interplay between past referral behavior and future referral value over time.

In a similar vein, *inactive_t* has a negative coefficient in M5 and M7, which suggests that users who have used the service recently make higher-value referrals. However, this result is also driven by unobserved heterogeneity. Specifically, users who have recently used the service are also more likely to know other potentially good customers (i.e., homophily-based matching). Indeed, Columns M6 and M8 show that fixed effects render the coefficient for *inactive_t* nonsignificant, confirming once again the importance of accounting for unobserved heterogeneity. Note, however, that M6 still suffers from confounding bias as it does not account for dynamic behavior. If users recommend the service first to their closest friends (and these friends are better prospects), then customers are likely to have already made their best referrals by the time they become experienced with the service. This would create a spurious negative association between experience level and referral value that should be attributed to the order in which referrals were made. It is only when we incorporate both fixed effects and dynamic behavior (as in M8) that we obtain a more accurate picture of how referral value evolves over time.

The positive coefficient for *prev_refs* and the negative coefficient for *prev_value* in M8 suggest that referral value increases as referrers make more referrals (Hypothesis 2c is supported), but less so if the value of prior referrals was high (Hypothesis 1d is supported). In terms of experience level, we find a positive and statistically significant coefficient for *rides_before*, implying that referral value increases as referrers become more experienced with the service (Hypothesis 1b is supported). Additionally, *rides* also has a positive and statistically significant coefficient, which implies that referral value is higher when referrers are using the service intensively (Hypothesis 1a is supported). Finally, *inactive_t* is not statistically significant, suggesting that long periods of inactivity do not affect referral value (conditional on the user making a referral).

In contrast to a logistic regression, the coefficients in a Poisson regression can be interpreted as relative marginal effects. For example, the coefficient of *rides_before* in M8 implies that when the experience level of a user increases by one ride (i.e., *rides_before* increases by one), then the value of future referrals made by that user increases by 0.28%. In Table 5, we report the relative marginal effects alongside the absolute marginal effects. The absolute effect in Table 5 is the average change in the number of rides made by referred accounts when the focal variable increases by one unit. For example, when the experience level of a user increases by one ride (i.e., *rides_before* increases by one), then the value of future referrals made by this user increases on average by 0.01 rides.

As mentioned when discussing the results on referral generation, marginal effects should be considered in the context of their cumulative impact. Thus, our model suggests that the value of referrals increases by 2.8% for every 10 rides accumulated by a user. In the case of the top 10% most experienced users in our sample, who completed more than 67 rides by the end of 2018, these results suggest that the value of their potential referrals could increase up to 18.76% ($67 \times 0.28\%$) relative to when they first started using the service.

The values in Table 5 suggest that past referral behavior is the most important factor in our analysis to characterize how referral value evolves throughout the customer's life cycle, but the relationship is intricate. The more referrals a user has made in the past, the higher the value of future referrals. Specifically, one additional past referral increases the value of the next referral by 3.88%. However, the value of future referrals decreases when the value of past referrals was high. Therefore, the overall impact of past referrals (value and quantity) on the value of future referrals may not be immediately clear from these results.

Nevertheless, note that the magnitude of the effect for *prev_refs* (3.88%) is small relative to *prev_value* (–5.20%). Thus, past referrals generally have a negative impact on the value of the next referral. For example, assuming that the first referral the user makes is equal to the average value (4.4 rides), these results imply that the value of subsequent referrals will decrease by 19%

Table 5. Average Marginal Effects for M8 in Table 4 (Referral Value)

	Absolute effect	Relative effect
rides_before	0.01 rides**	0.28%
rides	0.07 rides*	1.63%
prev_refs	0.17 rides*	3.88%
prev_value	−0.23 rides***	−5.20%

*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

on average ($3.88\% - 5.20\% \times 4.4$).[10] Only when customers have made a substantial number of referrals would past referral behavior contribute to higher-value referrals in the future. For instance, when a user has made six referrals of average value, the model estimates that the relative effect is positive: $3.88\% \times 6 - 5.20\% \times 4.4 = 0.4\%$. However, this occurred in very few cases in our sample.

6. Extended Empirical Analyses

In this section, we show the consistence of our findings under several robustness tests, provide additional details on the generalization of our findings, and extend our analysis to consider the evolution of referral attempts over time. For conciseness, we focus on testing the robustness of our results on referral generation. Robustness tests for referral value are presented in Appendix B.

6.1. Alternative Model Specifications

This section shows the robustness of our findings with respect to two potential concerns: the nonlinear specification of our model and confounding factors. The first concern relates to the fact that we must drop a large fraction of our data when estimating a logistic regression model with fixed effects. More specifically, we must drop the observations from all users who never made a referral because their individual-level fixed effects do not converge in the nonlinear specification. The second concern relates to unobserved time-varying factors that could confound our results, such as demand shocks driven by competing platforms. Table 6 considers four alternative model specifications to address these concerns, as described in more detail below. Our qualitative results are consistent across all these specifications, hence validating the robustness of our findings.

First, column M9 shows the results when estimating a linear probability model with fixed effects. Importantly, with this model, individual-level fixed effects converge even for individuals who do not make any referrals, so we can use the entire sample when estimating M9. Second, M10 shows the results when time fixed effects are incorporated at the weekly level in our main specification, allowing us to control for unobserved time shocks that may occur on a weekly basis. The same results also hold when we add time fixed effects at the biweekly level.

M11 incorporates additional controls associated with weekly demand for ride-sharing. These controls were built using more than 158 million trip records publicly released by the New York City Taxi & Limousine Commission (NYC TLC); these trip records are from both taxis and for-hire vehicles (e.g., ride-sharing platforms).[11] Controlling for demand shocks is important in our setting because they influence service usage (by definition) and could also potentially affect referral generation. We only use data from NYC (which is one of the cities in our data set) to estimate this model, which accounts for 78% of the observations in our sample, because these external data are only available for this city. These controls include:

Table 6. Alternative Models for Referral Generation

	M9 FEOLS	M10 Week effects	M11 Time controls	M12 2SLS[a]
rides_before	0.000102***	0.00212***	0.00339***	
	(10.11)	(6.21)	(8.09)	
rides	0.00284***	0.129***	0.135***	
	(34.60)	(43.27)	(38.21)	
inactive_t	−0.000578***	−0.105***	−0.101***	
	(−52.74)	(−68.61)	(−58.95)	
prev_refs	−0.0815***	−1.605***	−1.638***	
	(−86.88)	(−144.43)	(−127.25)	
prev_value	−0.00101*	−0.0456***	−0.0535***	
	(−9.27)	(−28.19)	(−25.90)	
log(*rides*)				1.667*
				(2.04)
N	9,527,831	918,749	719,785	208
Observation level	User-Week	User-Week	User-Week	Station-Week
Groups	400,000	30,917	24,042	4
Model type	OLS	Logit	Logit	2SLS
R^2 or pseudo-R^2	0.073	0.205	0.204	0.951
Time effects	Monthly	Weekly	Monthly	Time trend[b]
Extra time controls	No	No	Yes	No

Notes. All models include user/station fixed effects. Standard errors are adjusted by clustering at the user/station level. FEOLS, fixed effects ordinary least squares; 2SLS, two-stage least squares.

*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$. *t* Statistics in parentheses.

[a]The dependent variable is the logarithm of the number of referrals generated.

[b]A week time trend is used instead of weekly dummy variables due to the smaller data size.

weekly ride-sharing market share of our industry partner, weekly number of rides completed by our industry partner, weekly number of rides completed by other competing ride-sharing platforms, and weekly number of rides conducted by yellow and green taxis.

Finally, M12 incorporates weather conditions as instrumental variables to address potential confounding that is not addressed by M10 or M11. Specifically, users could undergo lifestyle changes at specific points in time that influence both their service usage and their referral generation (e.g., starting a new job, going back to school). Although we do not expect this behavior to be prevalent in our setting, this implies that our results for service usage could be upwardly biased due to factors that are not related to customer satisfaction. This type of potential confounding is not addressed in the other alternative specifications because it occurs at the user-week level (as opposed to just the user level or just the week level). Addressing this issue is particularly challenging in our context because randomized controlled experiments are infeasible: one cannot randomly assign a specific service usage to customers in order to assess the impact of usage on referrals. Therefore, we instead conduct an instrumental variable (IV) analysis that leverages exogenous variation induced by weather conditions as a shock to service usage.

Candidates for instrumental variables must satisfy two important requirements to be considered an exogenous shock (a valid instrument) in our setting. First, the relevance assumption must be met: the instrument must affect service usage. Second, the exclusion restriction must be met: the instrument must not affect referral generation by means other than its effect on service usage. This makes weather conditions a suitable instrument for our analysis. We expect weather conditions to affect service usage directly because users will tend to go out less when weather conditions are bad, whereas we do not expect weather conditions to affect referral generation by means other than an increase or decrease in service usage. We revisit the plausibility of these assumptions later, after discussing in more detail our IV specification.

The weather conditions we use as IVs are the same as the ones Aral and Nicolaides (2017) used to identify social contagion in exercise behavior. Specifically, we build instrumental variables using data from the National Oceanic Atmospheric Administration (NOAA) from four weather stations in NYC that tracked both temperature and precipitation during our observation period. We then use these weather conditions as instruments for the number of rides that were made close to these stations (service usage).[12] Finally, using a two-stage least squares regression (2SLS), we estimated the influence that the exogenously manipulated service usage had on the number of referrals made by individuals who took a ride close to the weather stations (referral generation). We used a log-log specification to account for the fact that referrals should be affected proportionally to the service usage experienced close to the weather stations. As M12 in Table 6 shows, a 1% increase in service usage due to weather conditions results in a 1.67% increase in referral generation, which is consistent with our findings at the user level.

Additionally, to increase the confidence in our results, we conducted multiple diagnostic tests for the IV analysis. First, we conducted a weak instruments test to evaluate the relevance assumption (*p*-value: 0.076) and a Sargan test to evaluate the over-identification restrictions (*p*-value: 0.201). Both tests corroborate the validity of our instruments at a 10% significance level. Based on the sensitivity approach proposed by Frank et al. (2013), we also found that potential violations of the exclusion restriction must be substantial to invalidate our results. The effect of weather conditions on referral generation must account for at least 19% of the estimated effect of service usage on referral generation for our results to stop being significant at a 10% level. Finally, we conducted a Wu-Hausman test to detect endogenous regressors (*p*-value: 0.581). This test implies that an ordinary least squares (OLS) regression is just as consistent as a 2SLS regression because there is no evidence of unobserved confounding between service usage and referral generation, which suggests that the IV analysis is unnecessary.

6.2. Subpopulation Analyses

This section shows how referral actions can evolve differently over time in different subpopulations of customers. Table 7 considers two types of analyses (described later) for our referral generation model. Appendix B includes similar analyses for our referral value model.

First, given that innovators (and early adopters) are often more influential and spread more WOM (Engel et al. 1969), our results could possibly vary depending on the stage of maturity of the platform. Assessing whether this is the case is important in the context of online platforms, given the critical role that referrals play in their growth. In our study, the user base was small compared with the overall population of the cities where the platform was operating, and as a result, there was substantial potential for the user base to grow further. This implies that our results may not apply to settings with saturated markets where the pool of potential new users is too small for referrals to effectively bring in new users. However, our data do include cities where the platform was operating at different stages of maturity, allowing us to assess whether there are systematic differences between how referral actions evolve for early adopters and late adopters.

In our next analysis, we create a new variable (*pioneer*) to segment users into early and nonearly adopters depending on the city in which they live in: *pioneer* = 1

Published in Information Systems Research on July 15, 2022 as DOI: 10.1287/isre.2022.1138.
This article has not been copyedited or formatted. The final version may differ from this version.

Table 7. Referral Generation Robustness Tests (for Different Cities and Riders' Segments)

	M13 Pioneer	M14 Referred
rides_before	0.00365***	0.00257***
	(8.82)	(5.44)
rides	0.135***	0.132***
	(38.31)	(29.23)
inactive_t	−0.100***	−0.101***
	(−59.47)	(−43.06)
prev_refs	−1.628***	−1.571***
	(−131.00)	(−96.13)
prev_value	−0.0534***	−0.0532***
	(−25.86)	(−20.00)
Interaction with	Pioneer	Referred
rides_before	−0.00403***	−0.00041
	(−6.06)	(−0.64)
rides	−0.0183**	−0.00217
	(−2.75)	(−0.36)
inactive_t	−0.0195***	−0.00401
	(−5.53)	(−1.39)
prev_refs	0.151***	−0.0291
	(7.28)	(−1.53)
prev_value	0.0218***	0.0128***
	(6.69)	(3.81)
N	918,749	918,749
Groups	30,917	30,917
Pseudo-R^2	0.205	0.204

Notes. All models include user and time fixed effects. Standard errors are adjusted by clustering at the user level.
$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. t Statistics in parentheses.

when the user adopted the service on a new market (i.e., the platform started its operations about a year ago), whereas *pioneer* = 0 when the user adopted the service on a more mature market (i.e., the platform has been operating for several years). We then estimate a new model (M13) that includes interaction terms with this new variable. Table 7 shows that M13 is consistent with our main findings (i.e., the sign and statistical significance of all coefficients remain the same), but the interaction terms are also statistically significant, which implies that the referral behavior of early adopters evolves differently over time.

These results show that early adopters are different from late adopters in at least three ways. First, the negative relationship between past and future referrals for early adopters is attenuated (*prev_refs* and *prev_value* have positive interaction coefficients). This is expected given that users are less likely to exhaust their pool of potential referrals when the service is new. Similarly, the relationship between referrals and service usage is also weaker (*rides_refs* and *rides* have negative interaction coefficients), possibly because intensive and accumulated service usage are not as representative of customer satisfaction and customer loyalty, respectively, for early adopters. Finally, the relationship between long inactive periods and less future referrals is stronger for early adopters (*inactive_t* has a negative interaction coefficient), which is consistent with studies that report that early adopters churn earlier (Lengyel et al. 2018).

For our second analysis, we estimate another model (M14) that includes interaction terms with the indicator variable *referred*, which has a value of 1 when the user was referred and a value of zero otherwise. We conducted this analysis because one plausible explanation for the positive impact of experience level on referral generation is that customers are more likely to become aware of the referral program as they become more experienced with the platform. If this is the case, then we would expect the interaction term with *rides_before* to have a negative coefficient because referred customers are aware of the referral program from the moment they start using the service (because they were referred themselves). However, our qualitative results remain the same as in M14, and this coefficient is not statistically significant, hence rejecting the hypothesis that the experience level relates to referral generation solely through an increased awareness about the referral program. Nevertheless, *prev_value* does have a positive and statistically significant coefficient. This is consistent with the fact that referred users are more likely to know people that are a good match for the firm (Van den Bulte et al. 2018), and as a result, these referrers are less likely to run out of potential people to refer.

6.3. Valuable Referral Generation

Another potential concern is the extent to which a subset of strategic customers is gaming the referral program to enjoy free rides (e.g., referred customers use the $10 free credit and then leave). To address this concern, we consider a logistic regression model that estimates the probability of making a referral *and* that the referred customer completes at least $k \geq 0$ rides. To the extent that such a gaming behavior is not substantial, our results should remain similar when increasing the value of k. In a way, this specification combines referral generation (Section 4.1) and referral value (Section 4.2) into a single model.

We use the same specification as in Equation (3) to conduct this robustness test, except that we only consider a referral to be successful (i.e., *ref* = 1) if the referred customer completes at least k rides during the first 13 weeks following account creation. Table 8 shows the results for k equals zero, one, two, and three (all referred customers who complete three rides or more are revenue-positive for the platform). When k = 0, this model is essentially the same as M4 in Table 2. As expected, the number of observations drops substantially as we increase k. This follows from the fact that the number of customers who made at least one successful referral decreases as k increases, and these are the only customers used to estimate our model (because of the user-level fixed effects). Nevertheless,

Table 8. Probability of Making Referrals with at Least k Rides ($k = 0, 1, 2, 3$)

	Rides0	Rides1	Rides2	Rides3
rides_before	0.00235***	0.00221***	0.00217***	0.00199***
	(6.92)	(5.69)	(4.78)	(3.82)
rides	0.130***	0.131***	0.132***	0.138***
	(43.60)	(38.72)	(32.61)	(29.63)
inactive_t	−0.104***	−0.103***	−0.104***	−0.100***
	(−67.97)	(−56.82)	(−41.96)	(−33.24)
prev_refs	−1.590***	−1.393***	−1.243***	−1.187***
	(−144.56)	(−119.20)	(−87.73)	(70.65)
prev_value	−0.0454***	−0.0951***	−0.159***	−0.204***
	(−28.11)	(−42.25)	(−50.54)	(−52.68)
N	918,749	724,678	503,149	385,341
Groups	30,917	24,301	16,528	12,517
Pseudo-R^2	0.204	0.201	0.207	0.222

Notes. All models include user and time fixed effects. Standard errors are adjusted by clustering at the referrer level.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$. t Statistics in parentheses.

the results are still statistically significant even when the sample size decreases. Moreover, the results are consistent with our previous findings. We find that the experience level has a positive effect on referral generation under various values of k. Similarly, current usage intensity (*rides*) and recency (*inactive_t*) remain significant for $k = 0, 1, 2, 3$.

Interestingly, the lagged variables also reflect the increasing importance of referral value relative to the probability of making a referral when increasing k: the coefficient for *prev_refs* decreases in magnitude and statistical significance, whereas the coefficient for *prev_value* increases in both magnitude and statistical significance even though the sample size shrinks. The former is the result of *prev_refs* having a negative coefficient in the referral generation model and a positive coefficient in the referral value model. The latter is the result of *prev_value* having a more prominent role for referral value than for referral generation.

6.4. Referral Attempts

A common limitation of referral data are that we only observe referral generation, but not failed referral attempts (i.e., attempts that did not translate into a successful referral). Although this is a common limitation when studying referral programs that are based on referral codes (Sun et al. 2021), there are two concerns that potentially arise from this limitation. The first concern relates to how referrals are generated: since users can share their referral code through various channels (e.g., social networks, email, blogs), looking at referral generation rather than referral attempts may only provide a partial picture when studying how referral behavior evolves with service usage, experience level, and past referral behavior. As mentioned before, we excluded from our data all users who made more than 10 referrals because our industry partner believes that such accounts may be connected to unsanctioned marketing activities. However, it is still possible that users who made less than 10 referrals are also involved in such activities. The second concern relates to the fact that we cannot infer whether the evolution of referral generation over time is driven by a change in the number of referral attempts or a change in the referral conversion rate.

Fortunately, our industry partner records the events when users rely on the ride-sharing app to share their referral code via one of the following channels: mobile text message, email, Facebook direct message, and Twitter direct message. Table 9 shows the percentage of referrers in our sample who used the app to share their code. Importantly, as the table shows, 83.8% of referrers used the app to share their referral code as opposed to share it outside the app (e.g., verbally during a conversation), so these data allow us to capture a large number of referral attempts.

Table 10 reports the estimation results for the models we used to address the above concerns. In column M15, we estimate our referral generation model using only the data from users who shared their referral code through the app in the same week or the week before making their referrals (we call this model "legitimate referrals"). We thus drop the observations from users who did not share their referral code via the app shortly before making one of their referrals. The intent behind this robustness test is to use only the data from users for whom we are confident that they only made legitimate referrals. Indeed, we would expect that users who shared their referral code in mass channels (e.g., online forums) to continue generating

Table 9. Percentage of Referrers Who Shared Their Code Through the App

Any channel	SMS	Email	Facebook	Twitter
83.8%	79.9%	19.4%	10.6%	6.4%

Table 10. Models That Incorporate Data from Referral Attempts

	M15 Legitimate referrals	M16 Referral attempts	M17 Referral conversions
rides_before	0.00210***	−0.00462***	0.00545***
	(4.41)	(−24.60)	(8.82)
rides	0.160***	0.0982***	0.0367***
	(39.98)	(57.25)	(6.23)
inactive_t	−0.098***	−0.0730***	−0.0370***
	(−49.91)	(−76.66)	(−8.99)
prev_refs	−2.175***	−0.575***	−1.541***
	(−119.83)	(−84.41)	(−70.39)
prev_value	−0.0447***	0.00289***	−0.0391***
	(−19.66)	(4.50)	(−14.08)
N	576,741	1,746,653	49,368
Groups	19,539	59,906	12,656
Dependent variable	Referral success	Referral attempt	Referral conversion
Pseudo-R^2	0.221	0.136	0.214

Notes. All models include user and time fixed effects. Standard errors are adjusted by clustering at the user level.
$*p < 0.05$; $**p < 0.01$; $***p < 0.001$. t Statistics in parentheses.

referrals without sharing their referral code via the app. All the results in this model are consistent with our previous findings.

In M16, we model the probability that a user will share her or his referral code through the ride-sharing app in a given week (i.e., referral attempts). We use the same specification as in Equation (3), except that the dependent variable is equal to one if the user made a referral attempt via the app in the focal week and zero otherwise. We use this model to assess whether our results are driven by referral attempts or by referral conversions. We find that the results for service usage and past referral behavior are consistent with our previous findings: namely, users share their referral code more often when they use the service intensively and frequently and less often when they have made past referrals. Thus, the relationship between these variables and referral generation is (at least partially) driven by a change in the number of referral attempts.

However, we find that the coefficient for experience level (*rides_before*) in M16 is negative, which is the opposite of what we found for referral generation. This suggests that users generate more referrals as they become more experienced *despite* making less referral attempts. While this empirical finding may seem counter-intuitive at first, it makes sense in light of the conflicting mechanisms we discussed when developing Hypothesis 1c. On the one hand, Garnefeld et al. (2013) argue that experience level should be negatively associated with referral generation because customers make less referral attempts as they become more experienced. On the other hand, Van den Bulte et al. (2018) argue that experience level should be positively associated with referral generation because the referral conversion rate increases as users become more experienced. M16 provides evidence that users indeed make less referral attempts as they become more experienced (as suggested by Garnefeld et al.), but the overall relationship between experience level and referral generation is positive because users become better at making referrals (i.e., their referral conversion rate improves) as they become more experienced (as suggested by Van den Bulte et al.).

In M17, we show additional support for the above claim by modeling how referral conversions evolve over time. Once again, we use the same specification as in Equation (3), except that the dependent variable is equal to one if the user made a referral attempt in the focal week *and* there was a successful referral that week, and zero otherwise. Furthermore, we only keep the observations from weeks when users made at least one referral attempt. As a result, the above dependent variable can be interpreted as a conversion conditional on the user making a referral attempt. As shown in M17, the experience level has a positive and statistically significant coefficient, implying that the conversion rate increases as users become more experienced. The coefficients for all other variables are also consistent with previous findings, which means that the relationship between these variables and referral generation are also partially driven by a change in the referral conversion rate.

7. Discussion and Conclusions

In this paper, we use a comprehensive panel data set to analyze how referral outcomes evolve throughout the customer's life cycle in the context of a ride-sharing platform. We find that users make more referrals (and of higher value) as they become more experienced with the service and when they use the service recently and intensively. In addition, as users make referrals, they become more likely to run out of friends to whom they can refer, hence leading to less referrals (and of lower value) in the future.

One important distinction between our paper and prior studies on referrals is the unit of analysis. While prior studies have focused on asking "which" customers make more and higher value referrals (i.e., a user-level analysis), we focus on the "when" (i.e., a time-user level analysis). This distinction is critical for the interpretation of our results. For instance, saying that experienced users make higher-value referrals (i.e., a "which" statement) is different from saying that users make higher-value referrals as they become more experienced (i.e., a "when" statement), and the former does not imply the latter or vice versa. For example, even if referral value does not increase with experience level, experienced users could be making higher-value referrals because they are more likely to have friends that would like the service. Namely, the statement "experienced users make higher-value referrals" could be driven by homophily in the referral process rather than by time-varying mechanisms, and it is not possible to uncover this finding based on a user-level analysis.

Our study makes a rigorous attempt to address this limitation, which is prevalent in the referral literature despite being fundamental for addressing endogeneity issues (Belo and Li 2018). More specifically, by incorporating user-level fixed effects in our models, we can account for unobserved (time-invariant) heterogeneity, such as homophily, and provide a more rigorous empirical assessment of how referral behavior evolves over time as a function of service usage and experience level. Importantly, our analysis refutes the negative causal link between experience level and referral generation that has been argued in previous studies based on cross-sectional data (Garnefeld et al. 2013). Furthermore, our results provide empirical support for other theoretically motivated mechanisms that positively link experience level with referral behavior.

Of course, user-level fixed effects do not account for time-varying, unobserved heterogeneity. This issue is particularly challenging in our context because randomized controlled experiments are infeasible: one cannot randomly assign a specific service usage to customers to assess the impact of usage behavior on referrals. Thus, in Section 6.1, we address time-varying confounding with various time-level controls and instrumental variables. This allows us to control for several sources of confounding factors that have been previously ignored, making this study an important step toward disentangling the relationship between service usage, experience level, and referrals.

Besides revealing the importance of accounting for these empirical challenges, our study also provides an alternative explanation for the empirical findings in Garnefeld et al. (2013). Specifically, our analysis in Section 6.4 is consistent with Garnefeld et al. (2013) in the sense that it suggests that users make less referral attempts as they become more experienced. Thus, the positive relationship between experience level and referral generation that we report in our study is driven by an increase in the referral conversion rate that outweighs the decrease in referral attempts.

In other settings, however, results may differ. Garnefeld et al. (2013) use data from a telecommunications provider, and such customers could be more likely to make referral attempts early on (explaining their results). Similarly, we find that past referrals generally imply less and lower value referrals in the future because users may run out of friends to refer, but this may not be the case for firms that allow non-unique customer referrals (e.g., in a retail context). Therefore, although we expect our results and insights to apply to other online platforms where customers frequently use the service and where making referrals requires a low effort (e.g., food delivery services, short-term lodging, fintech services), a promising avenue for future research is to assess the role of the context (e.g., type of service) as a moderating factor in the relationship between referral behavior and time-varying factors (e.g., service usage, experience level).

To our knowledge, our study is also the first to provide an empirical assessment of the relationship between past and future referral behavior. This relationship happens to be intricate. On the one hand, our results show that users become less likely to make new referrals after making high-value referrals, and the value of their next referrals decreases. On the other hand, our results (without fixed effects) also show that users who made high-value referrals in the past are more likely to do so again in the future. Importantly, the managerial implications of these two statements are quite different. If the manager's goal is to identify users who generate the most valuable referrals to allocate marketing resources, then the second statement implies that targeting customers who made high-value referrals in the past is a good strategy. However, if the manager wants to decide whether earlier referrals should be rewarded more than subsequent referrals, then the first statement implies that firms should be willing to invest more in earlier referrals.

This last statement suggests that referral programs could be potentially improved by tailoring them based on the evolution of customer behavior. We hope that our study will stimulate future research on how various design components of referral programs (e.g., call to action, rewards, referral request message) should evolve throughout customers' life cycle. For example, we show that customers make more referrals when they are using the service intensively, suggesting that timing calls to action and special offers (e.g., "Get an extra $X reward if you refer someone in the next X hours") according to service usage can be beneficial. Our results also suggest that firms may benefit from tailoring referral programs based on how referral value evolves throughout the customer's life cycle (e.g., by increasing referral rewards for more experienced users). Finally, another common design component for online platforms that benefit from network effects is the number of referrals that is required to unlock rewards (Belo and Li 2018). We find a decreasing return in the value of referrals as users make more referrals, which implies that firms should be cautious when increasing the number of referrals required to unlock rewards.

The results of our study also have managerial implications for other types of marketing efforts, such as advertising campaigns and customer retention programs. Our results suggest that short-term promotions and retention incentives based on service usage rewards may have benefits beyond increasing sales and reducing churn. More specifically, by increasing service usage, these marketing interventions may also encourage referrals. Thus, firms may also gain from the indirect benefit that their marketing interventions may have on referral behavior when making targeting decisions, even when the offers are not specifically designed to encourage referrals.

More broadly, we hope that our results will encourage firms to integrate their referral programs with other common types of marketing efforts that typically involve tracking customers over their life cycles. For instance, loyalty programs often involve rewards that are based on past purchases, and customer retention campaigns can be improved when designed around the stages that customers go before service cancellation (Ascarza and Hardie 2013). Because our results suggest that referral value increases as customers become more experienced, firms could effectively integrate referral and loyalty programs by allowing users to unlock higher referral rewards as part of loyalty programs. Similarly, referral outcomes and customer retention have been shown to be interrelated (Garnefeld et al. 2013, Van den Bulte et al. 2018), so jointly designing the rewards and the targeting decisions of referral and customer retention programs based on customer life cycle stages also offers an interesting direction for firms and future research to explore.

Acknowledgments

The authors thank our ride-sharing industry partner for sharing data and for insightful feedback and discussions; Zhitao Yin for suggesting Aral and Nicolaides (2017) as an example of how to use weather conditions as IVs; Aaron Lam for carrying out the data processing for the IV analysis; and Bill Greene and Baek Jung Kim for valuable input and feedback.

Appendix A. Analysis of Diminishing Marginal Effects

One potential concern about our main model specifications for referral generation and referral value is that they may not adequately model diminishing marginal effects. For example, the coefficients in a Poisson regression—such as the one we used to model referral value—imply a fixed percentage change in the dependent variable per unit change in the independent variables. This is inconsistent with our conceptual model because the marginal effect of satisfaction, loyalty, and knowledge about the firm's offering on referral generation and referral value should be decreasing. This appendix shows the results when we include quadratic terms in our main model specifications to address diminishing marginal effects. Overall, we find that both the statistical and economic significance of our findings increase under these alternative specifications.

Table A.1 shows the results when quadratic terms are included in the referral generation model (MA1) and in the referral value model (MA2). Importantly, the statistical significance of all our main results remains the same or increases when diminishing returns are included. Additionally, all quadratic terms in MA1 are statistically significant and have the opposite signs of the linear terms, which demonstrates the presence of decreasing marginal effects. In the case of MA2, only the quadratic terms for *prev_refs* and *prev_value* are statistically significant, but the statistical significance of all linear terms increases slightly

Table A.1. Main Specifications with Quadratic Terms

	MA1 Generation	MA2 Value
rides_before	0.00610***	0.00455***
	(9.60)	(3.70)
rides	0.209***	0.0284*
	(39.08)	(2.18)
inactive_t	−0.137***	0.0135
	(−48.73)	(1.22)
prev_refs	−2.47***	0.302***
	(−150.98)	(9.51)
prev_value	−0.0315***	−0.0808***
	(−15.82)	(−16.67)
$rides\_before^2$	−0.000017***	−0.000002
	(−8.45)	(−0.81)
$rides^2$	−0.00606***	−0.000945
	(−17.30)	(−1.41)
$inactive\_t^2$	0.00171***	−0.000556
	(18.20)	(−1.32)
$prev\_refs^2$	0.188***	−0.0392***
	(90.81)	(−7.55)
$prev\_value^2$	0.000133***	0.000307***
	(8.51)	(11.52)
N	918,749	39,516
Groups	30,917	13,160
Pseudo-R^2	0.226	0.527

Notes. All models include user and time fixed effects. Standard errors are adjusted by clustering at the user level.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$. *t* Statistics in parentheses.

Table A.2. Marginal Effects for MA1 in Table A.1 (Referral Generation)

	Average absolute effect	Average relative effect	Percentile 1% relative effect	Median relative effect	Percentile 99% relative effect
rides_before	0.02 pp***	0.55%	0.07%	0.58%	0.62%
rides	0.79 pp***	19.87%	4.62%	21.45%	22.37%
inactive_t	−0.48 pp***	−10.51%	−12.56%	−11.37%	−1.65%
prev_refs	−3.97 pp***	−80.20%	−89.63%	−84.81%	−3.42%
prev_value	−0.12 pp***	−2.90%	−3.12%	−3.01%	−1.28%

Note. The absolute effect is measured in percentage points (pp) whereas the relative effect is the percentage change in the probability of making a referral in a given week.
$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

when the quadratic terms are included. Therefore, diminishing marginal effects also play a significant role in the referral value model.

Tables A.2 and A.3 show the results of estimating marginal effects using the models in Table A.1. These estimates were obtained when including both linear and quadratic terms. With the exception of *prev_value* in referral generation (Table A.2), all these estimates are larger in magnitude than the ones reported in Tables 3 and 5. Therefore, the economic significance of our findings also increases when diminishing marginal effects are incorporated into our models.

Interestingly, Table A.3 shows that *prev_refs* is estimated to have a negative marginal effect on referral value for some individuals (see the 1%-percentile). This is consistent with the general result presented in the paper that past referral behavior is negatively associated with future referral value. This group of individuals corresponds to users who made a large number of referrals. Therefore, one explanation for this finding is that these referrers have already exhausted the pool of potentially good referral recipients in their social network.

Appendix B. Extended Analysis for Referral Value

This appendix includes similar analyses to the ones discussed in Sections 6.1 and 6.2, but for the referral value model. Models MB1, MB2, and MB3 in Table B.1 are analogous to the models presented in Table 6, and the models reported in Table B.2 are analogous to the ones from Table 7. Models MB4 and MB5 are meant to address potential overdispersion issues in the Poisson regression.

Models MB1, MB2, and MB3 are mostly consistent with our main findings, with two exceptions. First, the statistical significance level of *rides_before* drops to 10% when an OLS specification (MB1) is used. This is not a major concern because the qualitative interpretation of the result is the same, and the coefficient for *rides_before* in MB1 is similar to the absolute effect estimated in Table 5. Second, *prev_refs* is no longer statistical significant in MB3. One potential explanation is the substantial decrease in statistical power (recall that this model includes only NYC users). This is also not a major concern because MB3 is consistent with the main results: past referral behavior is negatively associated with future referral value.

MB4 and MB5 in Table B.1 are meant to address overdispersion. The concern here is that Poisson specifications assume that the conditional variance and the conditional expectation of the dependent variable are the same. This may result in incorrect estimates of the standard errors if overdispersion occurs, that is, if the data exhibits a greater variance than what is assumed by the Poisson specification.

We found that the overdispersion statistic (Person's chi-squared statistic) for our main Poisson model was quite high due to a small number of outliers with particularly large overdispersion values (20 of 39,516 observations). Specifically, the overdispersion statistic is 3,043.14 when the model is fit with these outliers and 2.11 when the model is fit without including them. MB4 reports the regression results when these outliers are omitted and the standard errors are scaled to correct for overdispersion. Overall, we find that some coefficients decrease in their statistical significance level when the standard errors are scaled, but they all remain statistically significant.

MB5 shows the results when a negative binomial regression is used instead of the Poisson specification. Although this alternative specification is subject to the incidental parameters problem (as discussed in Section 4.2), it is robust to overdispersion. As with MB4, some coefficients decrease in their statistical significance level in

Table A.3. Marginal Effects for MA2 in Table A.1 (Referral Value)

	Average absolute effect	Average relative effect	Percentile 1% relative effect	Median relative effect	Percentile 99% relative effect
rides_before	0.02 rides***	0.45%	0.39%	0.45%	0.46%
rides	0.10 rides*	2.39%	0.10%	2.59%	2.78%
prev_refs	0.82 rides***	17.18%	−24.84%	20.30%	30.12%
prev_value	−0.33 rides***	−7.55%	−7.74%	−7.68%	−5.58%

Note. The absolute effect is measured in number of rides in a 13-week period whereas the relative effect is the percentage change in the number of rides.
$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Table B.1. Alternative Specifications for the Referral Value Model

	MB1 FEOLS	MB2 Week effects	MB3 Time controls	MB4 Scaled standard error	MB5 Neg. binomial
rides_before	0.0119•	0.00285***	0.00257**	0.00286*	0.00121•
	(1.72)	(3.31)	(2.71)	(2.19)	(1.92)
rides	0.111*	0.0163*	0.0186*	0.0191•	0.0164**
	(2.45)	(2.21)	(2.14)	(1.75)	(3.11)
inactive_t	−0.00471	0.00142	0.00391	0.000961	0.00303
	(−0.30)	(0.28)	(0.69)	(0.13)	(0.84)
prev_refs	0.273***	0.0397**	0.0256	0.0546*	−0.00804
	(4.27)	(2.59)	(1.46)	(2.51)	(−0.77)
prev_value	−0.854***	−0.0537***	−0.0546***	−0.0621***	−0.0895***
	(−71.64)	(−12.24)	(−10.35)	(−11.00)	(−30.11)
N	39,516	39,516	29,366	39,496	39,496
Groups	13,160	13,160	9,757	13,160	13,160
Model type	OLS	Poisson	Poisson	Poisson	Neg. Binomial
R^2 or Pseudo-R^2	0.073	0.517	0.505	0.522	0.171
Time effects	Monthly	Weekly	Monthly	Monthly	Monthly
Extra time controls	No	No	Yes	No	No

Notes. All models include user-level fixed effects. Standard errors are adjusted by clustering at the user level.
•$p < 0.10$; *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$. *t* Statistics in parentheses.

MB5, but only *prev_ref* ceases to be statistically significant. As mentioned, this is not a major concern because MB5 is consistent with the main result that past referral behavior is negatively associated with future referral value.

As opposed to the findings discussed in Section 6.2 for the referral generation model, Table B.2 shows that most interaction terms are not statistically significant for the referral value model, with the exception of *prev_refs* in MB6. Therefore, the positive relationship between the number of past referrals and future referral value is mostly driven by early adopters. This is expected because referrers are likely to benefit more from past referrals when the service is new and when the pool of potential referral recipients is larger.

Table B.2. Subpopulation Analyses for the Referral Value Model

	MB6 Pioneer	MB7 Referred
rides_before	0.00230*	0.00304***
	(2.47)	(3.75)
rides	0.0191*	0.0241*
	(2.22)	(2.15)
inactive_t	0.00300	−0.00665
	(0.53)	(−0.78)
prev_refs	0.0205	0.0338
	(1.17)	(1.32)
prev_value	−0.0545***	−0.0511***
	(−10.40)	(−9.72)
Interaction with	Pioneer	Referred
rides_before	0.00038	−0.00042
	(0.25)	(−0.25)
rides	−0.00734	−0.0136
	(−0.46)	(−0.87)
inactive_t	−0.0111	0.00997
	(−0.92)	(0.97)
prev_refs	0.0643•	0.00814
	(1.89)	(0.27)
prev_value	0.00213	−0.00375
	(0.23)	(−0.46)
N	39,516	39,516
Groups	13,160	13,160
Pseudo-R^2	0.516	0.515

Notes. All models include user and time fixed effects. Standard errors are adjusted by clustering at the user level.
•$p < 0.10$; *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$. *t* Statistics in parentheses.

Endnotes

[1] See https://www.nytimes.com/2007/11/07/technology/07iht-07adco.8230630.html.

[2] See https://www.nielsen.com/us/en/insights/news/2013/under-the-influence-consumer-trust-in-advertising.html.

[3] We cannot disclose the name and details of the platform due to confidentiality.

[4] The variable *prev_value* takes a value of zero if the user has never made a referral in the past.

[5] Due to our nondisclosure agreement, we cannot report descriptive statistics.

[6] See https://cloud.r-project.org/web/packages/bife/vignettes/howto.html for a description of the software package we used.

[7] As with referral generation, our main results are qualitatively the same if we use weekly fixed effects instead.

[8] This extended analysis presents specifications that account for diminishing marginal effects in referral generation and referral value. Interestingly, the estimated average marginal effects are larger under these alternative specifications.

[9] All our results remain consistent when we include this group of outliers, except the coefficient for *inactive_t*. When we include this group of outliers, the coefficient for *inactive_t* is positive and statistically significant in M8.

[10] Recall that both *prev_refs* and *prev_value* have a value of zero when users have never made a referral in the past.

[11] Data source: https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page.

[12] The results we report correspond to rides made within a 10-mile radius. Rides that occurred in a 10-mile radius of these weather stations account for 78% of the rides we observe in NYC.

References

Allison PD, Williams R, Moral-Benito E (2017) Maximum likelihood for cross-lagged panel models with fixed effects. *Socius* 3: 2378023117710578.

Anderson EW (1998) Customer satisfaction and word of mouth. *J. Service Res.* 1(1):5–17.

Aral S, Nicolaides C (2017) Exercise contagion in a global social network. *Nature Comm.* 8(1):1–8.

Arellano M (2003) *Panel Data Econometrics* (Oxford University Press, Oxford, UK).

Arellano M, Hahn J (2007) Understanding bias in nonlinear panel models: Some recent developments. Blundell R, Newey W, Persson T, eds. *Advances in Econom. and Econometrics, Proc. 9th World Congress* (Cambridge University Press, Cambridge, UK), 381–409.

Ascarza E, Hardie BG (2013) A joint model of usage and churn in contractual settings. *Marketing Sci.* 32(4):570–590.

Bartolucci F, Nigro V (2012) Pseudo conditional maximum likelihood estimation of the dynamic logit model for binary panel data. *J. Econometrics* 170(1):102–116.

Belo R, Li T (2018) Referral programs for platform growth: Evidence from a randomized field experiment, Preprint, submitted August 15, https://dx.doi.org/10.2139/ssrn.3224330.

Berger J, Schwartz EM (2011) What drives immediate and ongoing word of mouth? *J. Marketing Res.* 48(5):869–880.

Biyalogorsky E, Gerstner E, Libai B (2001) Customer referral management: Optimal reward programs. *Marketing Sci.* 20(1):82–95.

Bolton RN, Lemon KN (1999) A dynamic model of customers' usage of services: Usage as an antecedent and consequence of satisfaction. *J. Marketing Res.* 36(2):171–186.

Brown JJ, Reingen PH (1987) Social ties and word-of-mouth referral behavior. *J. Consumer Res.* 14(3):350–362.

Buttle FA (1998) Word of mouth: Understanding and managing referral marketing. *J. Strategic Marketing* 6(3):241–254.

Castilla EJ (2005) Social networks and employee performance in a call center. *Amer. J. Sociol.* 110(5):1243–1283.

Chen MC, Chiu AL, Chang HH (2005) Mining changes in customer behavior in retail marketing. *Expert Systems Appl.* 28(4):773–781.

Dillahunt TR, Malone AR (2015) The promise of the sharing economy among disadvantaged communities. Inkpen K, Woo W, eds. *Proc. 33rd Annual ACM Conf. on Human Factors in Comput. Systems* (Association for Computing Machinery, New York), 2285–2294.

Engel JF, Kegerreis RJ, Blackwell RD (1969) Word-of-mouth communication by the innovator. *J. Marketing* 33(3):15–19.

Fernandez RM, Castilla EJ, Moore P (2000) Social capital at work: Networks and employment at a phone center. *Amer. J. Sociol.* 105(5):1288–1356.

Frank KA, Maroulis SJ, Duong MQ, Kelcey BM (2013) What would it take to change an inference? Using Rubin's causal model to interpret the robustness of causal inferences. *Edu. Evaluation Policy Anal.* 35(4):437–460.

Galenianos M (2013) Learning about match quality and the use of referrals. *Rev. Econom. Dynamics* 16(4):668–690.

Garnefeld I, Eggert A, Helm SV, Tax SS (2013) Growing existing customers' revenue streams through customer referral programs. *J. Marketing* 77(4):17–32.

Ghose A, Han SP (2011) An empirical analysis of user content generation and usage behavior on the mobile Internet. *Management Sci.* 57(9):1671–1691.

Greene W (2007) Fixed and random effects models for count data. Working paper, Stern School of Business, New York University, New York.

Hahn J, Newey W (2004) Jackknife and analytical bias reduction for nonlinear panel models. *Econometrica* 72(4):1295–1319.

Hallowell R (1996) The relationships of customer satisfaction, customer loyalty, and profitability: An empirical study. *Internat. J. of Service Indust. Management* 7(4):27–42.

Harrell FE (2015) Ordinal logistic regression. *Regression Modeling Strategies* (Springer, Berlin), 311–325.

Helm S (2003) Calculating the value of customers' referrals. *Management Service Quality.* 13(2):124–133.

Hong Y, Pavlou PA, Wang K, Shi N (2017) On the role of fairness and social distance in designing effective social referral systems. *Management Inform. Systems Quart.* 41(3):787–810.

Huang N, Burtch G, Gu B, Hong Y, Liang C, Wang K, Fu D, Yang B (2019) Motivating user-generated content with performance feedback: Evidence from randomized field experiments. *Management Sci.* 65(1):327–345.

Jing X, Xie J (2011) Group buying: A new mechanism for selling through social interactions. *Management Sci.* 57(8):1354–1372.

Jung J, Bapna R, Golden JM, Sun T (2020) Words matter! Toward a prosocial call-to-action for online referral: Evidence from two field experiments. *Inform. Systems Res.* 31(1):16–36.

Jung J, Bapna R, Gupta A, Sen S (2021) Impact of incentive mechanism in online referral programs: Evidence from randomized field experiments. *J. Management Inform. Systems* 38(1):59–81.

Koch O, Benlian A (2015) Designing viral promotional campaigns: How scarcity and social proof affect online referrals. Leidner D, Ross J, eds. *Proc. Internat. Conf. Inform. Systems* (Association for Information Systems, Atlanta), 656–676.

Kornish LJ, Li Q (2010) Optimal referral bonuses with asymmetric information: Firm-offered and interpersonal incentives. *Marketing Sci.* 29(1):108–121.

Krackhardt D, Nohria N, Eccles B (2003) The strength of strong ties. Cross R, Parker A, Sasson L, eds. *Networks in the Knowledge Economy* (Oxford University Press, Oxford, UK), 82.

Lancaster T (2000) The incidental parameter problem since 1948. *J. Econometrics* 95(2):391–413.

Lemon KN, Wangenheim FV (2009) The reinforcing effects of loyalty program partnerships and core service usage: A longitudinal analysis. *J. Service Res.* 11(4):357–370.

Lengyel B, Di Clemente R, Kertész J, González MC (2018) Spatial diffusion and churn of social media, Preprint, submitted April 4, https://arxiv.org/abs/1804.01349v1.

Neyman J, Scott EL (1948) Consistent estimates based on partially consistent observations. *Econometrica* 16(1):1–32.

Nielsen A (2012) Global trust in advertising and brand messages. Accessed February 12, 2020, https://www.nielsen.com/ssa/en/insights/report/2012/global-trust-in-advertising-and-brand-messages-2/.

Ryu G, Feick L (2007) A penny for your thoughts: Referral reward programs and referral likelihood. *J. Marketing* 71(1):84–94.

Schmitt P, Skiera B, Van den Bulte C (2011a) Referral programs and customer value. *J. Marketing* 75(1):46–59.

Schmitt P, Skiera B, Van den Bulte C (2011b) Why customer referrals can drive stunning profits. *Harvard Bus. Rev.* 89(6).

Stammann A, Heiß F, McFadden D (2016) Estimating fixed effects logit models with large panel data. Working paper, Heinrich-Heine-Universität Düsseldorf, Germany.

Sun T, Viswanathan S, Zheleva E (2021) Creating social contagion through firm-mediated message design: Evidence from a randomized field experiment. *Management Sci.* 67(2):808–827.

Sun T, Wei Y, Golden J (2019) Geographical pattern of online word-of-mouth: How offline environment influences online sharing. Preprint, submitted December 9, https://dx.doi.org/10.2139/ssrn.3497493.

Sun T, Viswanathan S, Huang N, Zheleva E (2021) Designing promotional incentive to embrace social sharing: Evidence from field and online experiments. *MIS Quart*. 45(2):789–820.

Van den Bulte C, Bayer E, Skiera B, Schmitt P (2018) How customer referral programs turn social capital into economic capital. *J. Marketing Res.* 55(1):132–146.

Wangenheim FV, Bayón T (2007) The chain from customer satisfaction via word-of-mouth referrals to new customer acquisition. *J. Acad. Marketing Sci.* 35(2):233–249.

Wei JT, Lin SY, Wu HH (2010) A review of the application of RFM model. *African J. Bus. Management* 4(19):4199–4206.

Wirtz J, Orsingher C, Chew P, Tambyah SK (2013) The role of metaperception on the effectiveness of referral reward programs. *J. Service Res.* 16(1):82–98.

Yi Y (1990) A critical review of consumer satisfaction. *Rev. Marketing* 4(1):68–123.

Zhang Y, Bradlow ET, Small DS (2015) Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Sci.* 34(2): 195–208.