

Dog Eat Dog: Balancing Network Effects and Differentiation in a Digital Platform Merger*

Chiara Farronato[†] Jessica Fong[‡] Andrey Fradkin[§]

January 6, 2022

Abstract

Mergers among digital platforms are increasingly receiving public and regulatory attention. These mergers may benefit users if network effects from a combined platform are large enough or may hurt users if the two platforms are differentiated and one of the platforms is shut down. We study the net effect of this trade-off in the context of the merger between the two largest platforms for pet-sitting services. We exploit geographic variation in pre-merger market shares and a difference-in-differences approach to causally estimate network effects. We find that users of the acquiring platform benefited from the merger because of network effects. However, users of the acquired platform were more likely to exit the market, for reasons including switching costs, coordination failures, and disintermediation. Network effects and attrition offset each other such that at the market level consumers are, on average, not substantially better off with a single combined platform than with two separate and competing platforms. Our results highlight the importance of platform differentiation even when platforms enjoy network effects, which has important implications for antitrust authorities and platform owners.

*We thank Ian Meeker and Hirotaka Miura for outstanding research assistance. We thank the company employees for sharing data and insights, and participants at INFORMS, the Marketplace Innovation Workshop, Marketing Science, the Platform Strategy Research Symposium, the Workshop in Information Systems and Economics, the Simons Institute, the QME conference, and the NBER Summer Institute for valuable comments. We also thank seminar participants at Boston University, Harvard University, University of Maryland, University of Michigan, Toulouse School of Economics, the Bank of Italy, DIW Berlin, the Online Economics of Platforms Seminar, Louvain and Paris Economics of Digitization Seminar, the NYU Strategy Economics Workshop, University of Chicago, CU Boulder, and the CEPR Virtual IO seminar for very helpful feedback. Rover reviewed the paper to make sure that required confidential information was reported accurately. None of the authors have any material financial relationship with entities related to this research.

[†]Harvard Business School and NBER, cfarronato@hbs.edu

[‡]University of Michigan, jyfong@umich.edu

[§]Boston University, fradkin@bu.edu

1 Introduction

Recent mergers and acquisitions of digital platforms have increased public and policy attention to the dominance of such platforms.¹ Two competing narratives have emerged in this debate. On the one hand, users can benefit when two platforms combine (e.g. Microsoft and Hotmail) if there are strong network effects, in which the value per user increases with the number of other users on the platform. On the other hand, users may be better off when the acquired and acquiring platform remain separate (e.g. Expedia and Orbitz) if the platforms are horizontally differentiated and users have varying preferences. We study the relative importance of these forces in a merger between two digital platforms competing in the local services industry, in which the largest platform acquired and combined with its largest competitor.

We find that, on average, users are not significantly better off with a single platform compared to two competitors. This is true despite significant network benefits experienced by the acquiring platform and no changes in platform prices. The network effects are offset by the fact that many users of the acquired platform chose not to switch to the acquiring platform. Since the acquired platform was shut down, acquired users who strongly preferred their original platform were actually hurt by the merger. At the market level, this heterogeneity in user preferences across platforms counterbalances platform-level network effects.

A key challenge that our study is able to overcome is the measurement of network effects, which are one of the defining characteristics of platforms (Rochet and Tirole (2003) and Cusumano et al. (2019)). In the specific context of online marketplaces, network effects may increase the level or quality of platform-intermediated exchanges following an increase in the number of users. But network effects are difficult to quantify because platform growth is typically endogenous.

We have the unique opportunity to measure network effects from the combination of two online marketplaces for pet-sitting services. In March 2017, Rover, the largest US pet-sitting platform, acquired DogVacay, their closest and largest competitor. A single

¹Examples include Microsoft acquiring Hotmail in 2012, Facebook acquiring Instagram in 2012, Expedia acquiring Orbitz in 2015, and Match.com acquiring Hinge in 2019.

platform emerged from this acquisition since DogVacay was shut down within 4 months of the acquisition announcement. The setting is unique in that we observe data from *both* platforms before and after the acquisition, and we are able to identify the same users across the two platforms.

This acquisition provides an excellent natural experiment for 1) testing for the presence of network effects, and 2) evaluating whether network effects are large enough to justify a single combined platform over maintaining two separate and competing platforms. First, the local nature of services exchanged means that interactions in one city do not affect interactions in another city, so we can treat each geography as a separate market (Cullen and Farronato (2021)). Second, the two platforms were similar in the services exchanged and the way in which buyers search for service providers. These similarities imply that the potential for network effects to arise is high. Third, prior to the acquisition, the two platforms varied in their market shares across cities, which means that some cities experienced bigger increases in the number of users interacting with one another compared to other cities. Finally, the acquiring platform did not increase its nominal or actual commission fees, a potential confound that may offset the benefits of the acquisition to platform users.

The presence and size of network effects generated by combining the two platforms depends on the level of competition before the acquisition. We show that the two platforms are comparable in size and they are active in the same geographies. We also show that multi-homing is limited, thus reducing pre-merger interactions of buyers on one platform with sellers on the other. Finally, we find evidence that providers charge similar prices when selling on both platforms, suggesting that the two platforms are indeed similar in the eyes of some users. These preliminary analyses show that combining the two platforms does affect the number of people with whom each user can interact, implying that there is scope for network effects to arise.

There are several ways in which the presence and behavior of one user can affect the utility of other users in platforms like ours, which are essentially online marketplaces where many buyers and sellers exchange goods or services. First, more buyers can increase the profits of sellers through increased demand, and more sellers can improve the outcomes of

buyers by providing better matches and prices.² These spillovers imply that the number of buyers *relative* to sellers affects the surplus created by the platform and how it is distributed across users. Second, and the specific focus of this paper, a change in the *absolute* number of buyers and sellers holding constant their relative shares may make buyers and sellers better off due to network effects. This may occur, for example, if greater variety on both the demand and the supply side results in more and higher quality matches.

Our first question is whether network effects even exist in platforms like ours. To answer this question, we study the *effect of the merger on the acquiring platform*, exploiting variation in pre-merger market shares that are at least in part explained by differences in the growth strategies of the two platforms. With network effects, one would expect the acquiring platform to benefit more in geographies receiving a bigger influx of users from the acquired platform.

Our second question is whether network effects are large enough to justify a single platform over two. This is both a managerial and policy-relevant question. From a managerial perspective, the acquiring firm has the opportunity to continue operating the two platforms separately, or to shut down one platform and invite its users to join the other. From a policy perspective, the antitrust regulator has the authority allow or block a merger altogether, or even to stipulate that the acquiring firm continue operating both platforms separately. To answer this question we study the *effects of the merger on the market*, aggregating data from both platforms. If network effects are large enough, combining the two platforms would lead to larger benefits in geographies where each platform had 50% of the market before the merger compared to geographies where one platform was already dominant before the merger. This is because network effects imply that the benefits of doubling platform users for all users are bigger than moving a small number of users to an already large platform.

We use a difference-in-differences strategy to measure the effect of merging the two platforms, comparing outcomes before and after the acquisition, and across geographies with different market shares. We explicitly address selection into market shares and spillovers between geographies, which may result in bias if left unaddressed. Specifically, we match

²Some of these effects, namely changes in price as a function of aggregate demand and supply, are purely competitive effects, but other externalities across buyers and sellers are often called cross-side or indirect network effects.

geographies where Rover, the acquiring platform, was not dominant before the acquisition — our “treated” units — to geographies where Rover had more than 80% market share. We match these geographies based on their pre-acquisition number of active sellers across the two platforms. To address spillovers across geographies, in robustness checks presented in the Appendix, we use market definitions that are coarser than zip codes and are based on users’ search behavior.

We find evidence of network effects at the *platform* level. In particular, we find that existing users on Rover increased their usage of the platform post-merger. They were more likely to submit requests, which resulted in more transactions given that request match rates for those existing users stayed relatively constant. However, even if existing DogVacay users benefited from network size, they decreased their platform usage after the merger *relative to* existing Rover users. In fact, DogVacay users were more likely to exit the market post-merger. Many of these users chose not to migrate their profile to Rover, and those who migrated transacted less frequently and matched at lower rates than comparable Rover users. The attrition almost perfectly offsets the increased usage of Rover users so that at the market-level we find no evidence that the combined platform substantially improves market outcomes more than the sum of the two separate platforms: not on the extensive margins such as user adoption, retention or total transactions, nor on the intensive margins, such as match rates or ratings.

Our results imply that even if network effects are strong in online platforms, preference heterogeneity can offset the benefits of a single platform compared to multiple competing platforms. This result is true across different types of geographies: geographies with a small versus large baseline number of users, and geographies where users have lower versus higher propensity to multi-home.

By revealed preferences, DogVacay users prefer the DogVacay platform to the Rover platform. This may be due to a variety of factors, including the user interface, search algorithms, and composition of DogVacay users. Indeed, we find that repeat transactions help explain DogVacay user attrition. In particular, DogVacay users who have multiple transactions with the same trading partner prior to the merger are more likely to leave the market relative to similar Rover users. This effect is even larger when the prior transaction

partner does not migrate to Rover. These results point to two related mechanisms. The first is disintermediation, where users transact off the platform. The second is a coordination failure, where users can't find each other on the new platform because of the platform interface or because search algorithms are not tailored to their specific preferences.

The attrition results inform the generalizability of our main finding, that platform differentiation offsets network effects. Both disintermediation and coordination failures are more likely to be present post-merger in platforms where repeat interactions are relatively important, such as child or elderly care, as opposed to platforms where matches are unlikely to occur with the same trading partner, such as ride-sharing.

The rest of the paper is structured as follows. In Section 2, we present the relevant literature. Section 3 presents a stylized model motivating our empirical analysis. Section 4 describes the data and the natural experiment while Section 5 presents our empirical specification. Results are in Section 6. In Section 7, we conclude by discussing implications for platform strategy and antitrust regulation.

2 Literature Review

In this section, we present the mostly theoretical literature on platforms and network effects, and describe how the setting in this paper is ideal for studying network effects empirically.

Early theoretical work focuses on competition in the presence of network externalities and product compatibility (Katz et al. (1985) and Farrell and Saloner (1985)), but the pioneering models of multi-sided platforms came with Rochet and Tirole (2003), Caillaud and Jullien (2003), Parker and Van Alstyne (2005), and Armstrong (2006), which were later generalized by Weyl (2010). In their models, two characteristics define platform businesses. The first characteristic is that platforms attract multiple user groups and enable interactions between them – e.g. buyers and sellers, or advertisers and social media users. The second characteristic is the presence of positive network effects, which imply that surplus per user is an increasing function of the number of participating users. These models typically focus on cross-side network effects, where each user (e.g. buyer) is directly affected by the number of users in other groups (e.g. sellers). The focus of these early models was to study platform

pricing strategies to attract multiple user groups. Other strategic choices, such as entry, vertical integration, and degree of openness have been studied by Zhu and Iansiti (2012), Hagiu and Wright (2014), and Boudreau (2010), among others. A crucial implication of this theoretical literature is that because of network effects, the value per user increases in the number of platform users. Two other theoretical papers, Tan and Zhou (2020) and Nikzad (2020), study how network effects, product variety, and prices lead to competition among platforms having ambiguous effects on consumer surplus. Our work adds an empirical focus to this literature, by estimating whether user outcomes improve with the number of other participating users on a platform and whether these network benefits more than offset reduced product differentiation. Our results on the lack of market-level improvements emphasize the importance of product variety in counterbalancing network benefits.

Another related stream of theoretical literature on platforms focuses on multi-homing, i.e. the propensity of users to join and use multiple substitute platforms. A couple of papers look at multi-homing users on both sides of the interaction (Caillaud and Jullien (2003) and Bakos and Halaburda (2019)), while most papers either assume single-homing or allow for multi-homing by only one side of users. When multi-homing is limited to at most one side, the strategic interdependence between the two sides implies that a platform may maximize profits by subsidizing one side to charge the other (Weyl (2010)). We contribute to this literature by providing empirical evidence on the extent of multi-homing in practice, finding that multi-homing, albeit somewhat limited, is predominantly concentrated on the seller side.

The empirical literature on network effects dates back to Greenstein (1993), Gandal (1994), and Saloner and Shepard (1995), who show early evidence that network effects are present in federal computer procurement, in the adoption of computer spreadsheet programs, and in banks' adoption of ATMs, respectively. One of the first to empirically study and find evidence of positive cross-side network externalities is Rysman (2004) in the market for Yellow Pages.

Data on how users interact with each other on digital platforms have allowed recent studies to estimate a particular manifestation of network effects, i.e. how the number of matches between the two sides of users changes as a function of aggregate user participation.

In the market for domestic tasks and errands, Cullen and Farronato (2021) do not find evidence of increasing returns to scale in matching. Analogous findings were confirmed in home sharing by Fradkin (2018) and Li and Netessine (2020), and in online dating by Fong (2019). Kabra et al. (2017), on the other hand, find positive returns to scale in ride-sharing. Reshef (2019) studies how new sellers on a platform affect established sellers using data from the Yelp delivery platform and Grubhub. Chu and Manchanda (2016) find cross-side network effects on Taobao.

Similar to our paper, Li and Netessine (2020) study a migration of listings onto a short-term rental platform to study how the increase in product choice affects match rates. They find a substantial decrease in match rates, which is attributed to search frictions. Our study complements this existing work because we are able to measure user behavior on the acquired platform prior to the acquisition. This allows us to characterize the effects of merging two platforms not only at the level of the acquiring platform, but also at the market level, accounting for differences in user composition. Our data also allows us to understand the role of multi-homing, which to our knowledge has never been possible before in the digital setting.³ Finally, the ability to measure network effects on multiple dimensions, from the extensive margins – i.e. number of transactions – to the intensive margins – i.e. match rates and match quality, provides a comprehensive analysis of the effects of the merger at every stage of the transaction funnel.

In addition to measuring network effects, our work relates to the literature on digital platform competition. Many papers have studied platform competition but mostly in a non-digital setting or focusing on competition between digital platforms and traditional service providers (Seamans and Zhu (2014), Farronato and Fradkin (Forthcoming), and Lam et al. (2021) among others). Rysman (2004) finds that platform competition is better for user surplus due to lower market power, although Chandra and Collard-Wexler (2009) find that concentration in the Canadian newspaper industry did not lead to higher prices for either newspaper subscribers or advertisers. Gowrisankaran and Stavins (2004) study banks' adoption of automated clearinghouse (ACH) electronic payment systems, Berry and Waldfogel (1999) and Jeziorski (2014a,b) study radio stations, and Tucker (2008) study the

³The exception is the known consumer and merchant behavior over credit card use, where consumers own and merchants accept multiple credit cards. See Bakos and Halaburda (2019) for survey evidence.

adoption of a video-messaging technology in a bank. Dubé et al. (2010) study market tipping and find that network effects can lead to a strong increase in concentration in the market for video game consoles. More recently Kawaguchi et al. (2020) conduct simulated merger analysis of mobile apps while He et al. (2021) attribute price increases on Amazon following the demise of Toys R Us to the platform’s market power. Unlike our work, these papers often focus only on the extensive margins of user participation ignoring usage intensity and match quality.

3 Theoretical Framework

This section presents a simple model of online matching platforms that includes network effects and platform differentiation. This model gives us simple expressions for users’ utilities and will guide our analysis of the platform merger. We focus on buyers, but a symmetric analysis can be done for sellers.

Our model is a Hotelling model with network effects. We define a market to be the activity of buyers and sellers within a local geography – e.g., a zip code – and short time period – e.g., a month. In a market there is a unit mass of sellers and a unit mass of buyers distributed uniformly on the (0,1) line. We assume that the number of buyers relative to sellers is fixed – equal to 1 for simplicity – and that buyers and sellers’ utilities are such that competing platforms have the same ratio of participating buyers relative to sellers.

There are two platforms: platform 1, the acquiring platform, and platform 2, the acquired platform. Platform 1 is located at 0 while platform 2 is located at 1. A buyer i located at point d_i has utility for platform 1 equal to $u_{i1}(s) = v(s) - d_i$, where s is the share of users using platform 1. Assuming that the market is always covered, i.e. nobody chooses the outside option, the share of users on platform 2 is $1 - s$ and the corresponding utility for i is $u_{i2}(s) = v(1 - s) - (1 - d_i)$. *Network effects* exist whenever $v(\cdot)$ is increasing in its argument. Horizontal preferences are given by the parameter d_i .

Consumers choose which platform to join as a function of platform utilities and a joining cost. We normalize platform 1’s joining cost to 0, and we denote platform 2’s joining cost c , which can be positive or negative. This means that user i joins platform 1 if and only

if $u_{i1}(s) \geq u_{i2}(s) - c$. Importantly, c is only present to rationalize different market shares across multiple geographies, but does not affect the utility obtained on the platforms after joining. Therefore, we will ignore c when computing the benefits of competition relative to a monopoly.

The equilibrium share of users joining platform 1 is given by the following indifference condition for the marginal user: $v(s) - s = v(1 - s) - (1 - s) - c$. Depending on $v(\cdot)$, we can have a single or multiple equilibria. The extremes, $s = 0$ and $s = 1$ can be equilibria. $s = 1$ is an equilibrium whenever $v(1) - 1 > v(0) - c$, i.e., the utility of users located at 0 from joining the farther platform if everybody else is on it is greater than being alone on the closest platform. Symmetrically, $s = 0$ is an additional equilibrium whenever $v(0) < v(1) - 1 - c$. For our purposes, equilibrium selection does not matter. Given an equilibrium s , the average utility provided by platform 1 is $\bar{u}_1(s) = v(s) - 0.5s$, the average utility provided by platform 2 is $\bar{u}_2(s) = v(1 - s) - 0.5(1 - s)$, and the average utility in the market is the weighted average of utilities from the two platforms, where the weights depend on the market shares: $s\bar{u}_1(s) + (1 - s)\bar{u}_2(s)$.

When platform 1 acquires platform 2 and users from platform 2 are required to migrate to platform 1, the average utility provided by platform 1 is $\bar{u}_1(1) = v(1) - 0.5$, which also coincides with the market-level average utility.⁴ Figure 1a plots the average utility for users of platform 1 as a function of s . The figure shows how the component of utility that depends on platform size $v(s)$ is increasing in s , while the component of utility that depends on the distance from platform 1 is negative and decreasing in s .

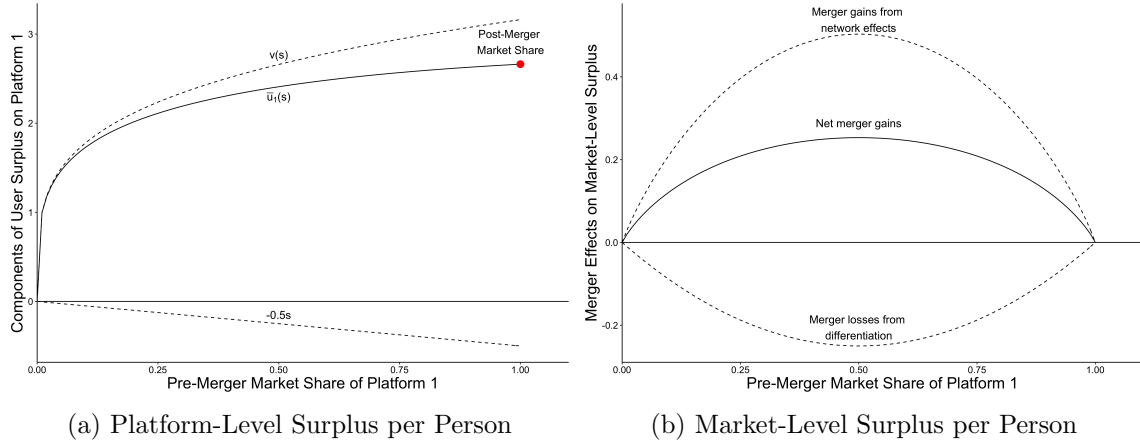
To compare how utilities change after the merger of the two platforms, we start with users who were on platform 1 before the merger. Their average utility changes as follows:

$$\bar{u}_1^1(1, s) - \bar{u}_1(s) = v(1) - v(s),$$

where $\bar{u}_1^1(1, s) = v(1) - 0.5s$ denotes platform 1's utility after the merger to users who were on platform 1 before the merger. Network effects imply that the increase in utility for

⁴Note that platform 1 could acquire platform 2 but could keep both platforms operational, such as in the case of Zillow and Trulia. Our theory is not designed to explain the effects of just changes in ownership structure.

Figure 1: Consumer Surplus at Platform and Market Level



The left panel plots surplus per user of users in platform 1, in aggregate and split into its two components. Average utility on platform 1 is $\bar{u}_i(s) = v(s) - 0.5s$. The solid line plots $\bar{u}_i(s)$, the dashed lines plot $v(s)$ – component of utility that depends on platform size – and $-0.5s$ – component of utility that depends on the distance from the platform. For the plot we assume $v(s) = (100s)^{0.25}$. Utility from platform 2 is exactly symmetric around 0.5. The right panel plots the change in market-level surplus after the merger as a function of platform 1’s pre-merger market share. The post-merger utility per user is $\bar{u}_1(1) = v(1) - 0.5$, while the pre-merger utility is $s\bar{u}_1(s) + (1-s)\bar{u}_2(s) = sv(s) + (1-s)v(s) - (0.5s^2 + 0.5(1-s)^2)$. The positive dashed line plots $v(1) - [sv(s) + (1-s)v(s)]$, which represents the merger gains from network effects. The negative dashed line plots $-s(1-s)$, which represents the merger losses from reduction in platform differentiation. The solid line combines the two by plotting $\bar{u}_1(1) - [s\bar{u}_1(s) + (1-s)\bar{u}_2(s)]$.

existing users of platform 1 is positive for any $s < 1$ and is decreasing in s :

$$\bar{u}_1^1(1, s) - \bar{u}_1(s) \text{ is decreasing in } s. \quad (1)$$

In words, the equation states that with network effects, the increase in average surplus for existing users of platform 1 is bigger in markets where platform 1 was smaller before the merger. That happens because the acquiring platform receives a bigger influx of users from platform 2 when s is small compared to when s is large.⁵

To evaluate the role of horizontal preferences, we compare the post- and pre-merger

⁵Note that comparing the average utility of all users in platform 1 after the merger with the average utility of all users in platform 1 before the merger cannot separate network effects from selection effects. Specifically, $\bar{u}_1(1) - \bar{u}_1(s) = [\bar{u}_1^1(1, s) - \bar{u}_1(s)] - 0.5(1-s)$. The first term in square brackets, $\bar{u}_1^1(1, s) - \bar{u}_1(s)$, represents the change in utility for the users who were already using platform 1 before the merger. The second term, $0.5(1-s)$, represents the reduction in utility coming from the fact that the users migrating from platform 2 are located farther away from platform 1 than the existing users and thus have higher average travel costs.

utility of users who chose platform 2 before the merger:

$$\bar{u}_1^2(1, s) - \bar{u}_2(s) = v(1) - v(1 - s) - s, \quad (2)$$

where $\bar{u}_1^2(1, s) = v(1) - 0.5(1 + s)$ is platform 1's utility to users who were on platform 2 before the merger. Like before, the first component, $v(1) - v(1 - s)$, is the increase in utility due to network effects, which is positive and increasing in s . The second component is the increase in travel costs due to the migration to platform 1.

The predictions of how utility changes after the merger for users of platform 2 are ambiguous. Depending on whether network effects dominate over horizontal preferences, users of platform 2 may be better off or worse off after the merger. We note however, that there is a close relationship between the gains of users from platform 1 and platform 2. In particular, consider the change in utility for platform 1's users in a market where platform 1 has market share s and the change in utility for platform 2's users in a market where platform 1 has market share $1 - s$:

$$[\bar{u}_1^2(1, 1 - s) - \bar{u}_2(1 - s)] - [\bar{u}_1^1(1, s) - \bar{u}_1(s)] = s - 1.$$

The merger gains for these two set of users are the same except for the increase in travel costs for platform 2 users. Given that s can only take values between 0 and 1,

$$[\bar{u}_1^2(1, 1 - s) - \bar{u}_2(1 - s)] - [\bar{u}_1^1(1, s) - \bar{u}_1(s)] \text{ is negative and increasing in } s. \quad (3)$$

Are network effects large enough that a single platform can create more value for users than two competing platforms? For this to be true, network effects need to dominate over horizontal preferences. The change in market-level average utility from the merger is as follows:

$$\bar{u}(1) - [s\bar{u}_1(s) + (1 - s)\bar{u}_2(s)] = [v(1) - sv(s) - (1 - s)v(1 - s)] - s(1 - s).$$

The first term, $v(1) - sv(s) - (1 - s)v(1 - s)$, represents the gains from scale. The second term, $s(1 - s)$, represents the losses from the reduction in platform differentiation. Figure 1b

plots the two terms separately as dashed lines. Whether a single platform or two competing platforms create more user value depends on which of the two terms dominates. If network effects dominate (as shown in Figure 1b), we have that $\bar{u}(1) - [s\bar{u}_1(s) + (1-s)\bar{u}_2(s)]$ is positive and reaches a maximum at $s = 0.5$:

$$\bar{u}(1) - [s\bar{u}_1(s) + (1-s)\bar{u}_2(s)] \text{ is decreasing in } |s - 0.5|. \quad (4)$$

Our theory provides testable implications of network effects realized in a merger and shows how the benefits of the merger depend on the relative importance of network effects and differentiation. We discuss two extensions to our model in Appendix A. In particular, we allow a share of platform 2 users to choose the outside option and not migrate to platform 1 at all post the shut-down of platform 2. This modification helps to explain some of our empirical results. We also extend the model to allow for heterogeneity in market size and the degree of multi-homing. In the rest of the paper, we conduct tests of the model and measure whether network effects overcome platform differentiation.

4 Setting and Data

We have proprietary data from “A Place for Rover, Inc.” (Rover). Founded in 2012 in Seattle, Rover was the largest online platform for pet care services in the US, with a valuation of \$970 million as of 2018.⁶ At the time, Rover processed roughly one million bookings per month. DogVacay was a nearly identical platform, founded in 2012 in Santa Monica.

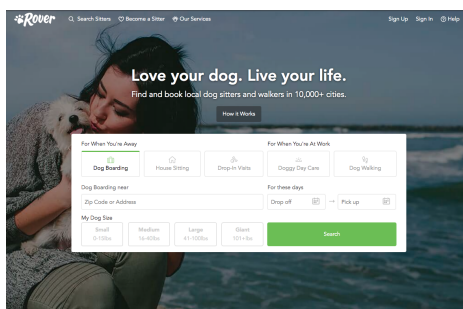
The pet industry market is large and growing. According to the American Pet Products Association,⁷ in 2019 pet owners in the US spent \$95.7 billion on their pets, including \$10.7 billion in services like boarding, grooming, training, pet sitting, and walking. That constitutes a 5.5% increase over the previous year. In the US, 84.9 million households, or 68% of all households, own a pet. Of them, 75% own a dog.

Dog owners (buyers) use Rover – and DogVacay before the acquisition – to find pet

⁶<https://www.wsj.com/articles/rover-raises-125-million-as-dog-sitting-war-heats-up-1527166801> (accessed July 2019).

⁷<https://www.americanpetproducts.org/pr> (accessed April 2020).

Figure 2: Rover’s and DogVacay’s Landing Pages



(a) Rover.com, March 2017.



(b) Dogvacay.com, March 2017.

The figures show the landing page of Rover and DogVacay before the acquisition. The screenshots are accessible on Wayback Machine (<https://web.archive.org/web/20170307101746/https://www.rover.com/> and <https://web.archive.org/web/20170228165616/https://dogvacay.com/>)

care services from sitters (sellers).⁸ The services range from dog walking to in-home pet grooming, but their largest category is dog boarding. Before the acquisition, Rover and DogVacay were the largest players in the online dog boarding market. At the time, the next largest competitor was Wag Labs (Wag). Wag, which mainly offered dog-walking services, started offering overnight boarding only in 2016,⁹ although it never grew to become their largest service category. In 2017, Rover earned five times higher revenues than Wag.¹⁰ Offline competitors include more traditional businesses like kennels and dog hotels, and more informal alternatives such as friends and family.

On the surface, Rover and DogVacay appear to be close substitutes, especially in comparison to competing platforms in other industries. In particular, Rover and DogVacay had similar interfaces (Figure 2) and transaction flows, which remained constant at least until the end of our study. When buyers need pet care services, they initiate a search for sellers available in the preferred category,¹¹ for a given location, and for the dates needed. As is typical in online platforms for local services, buyers then see a list of search results for

⁸It is fairly easy to join the platform as a pet sitter. One of us signed up on Rover by creating a sitter profile. Platform approval was quickly granted after a general background check. Additional background checks can be performed at the sitter’s will (<https://www.rover.com/background-checks/>, accessed July 2020).

⁹<https://www.vox.com/the-goods/2018/9/12/17831948/rover-wag-dog-walking-app>, accessed December 2020.

¹⁰<https://secondmeasure.com/datapoints/wag-rover-dog-walking-sales/>, accessed December 2020. Note that this figure includes total sales, not just from dog boarding.

¹¹The service categories include pet overnight boarding, sitting, drop-ins, daycare, and walking.

available sellers ranked by the companies' proprietary algorithms. For each seller displayed in search results, buyers see their name, picture, location, online ratings, and nightly price. Buyers can then choose to contact sellers to discuss their needs and confirm availability. An exchange is not finalized until both users accept the transaction. Transactions come with reservation protection, trust and safety support, and a secure payment system.

A deeper comparison uncovers a number of differences between the platforms. Platforms use proprietary algorithms to rank sitters in search results, weighing sitter characteristics differently.¹² DogVacay used to offer a 'meet and greet' option before finalizing a match whereas Rover did not. Lastly, sorting across the platforms could create differences in the user experience, either due to path-dependence or due to strategic decisions by the platforms regarding which types of users to attract (Halaburda et al. (2018)).

Just before the acquisition, both Rover and DogVacay took about 20% of gross transaction volume in commission fees, up from 15% when they first started. Sellers would set the prices for their services.¹³ As of 2018, fees are divided into a provider (seller) fee and a owner (buyer) fee. The provider fee is 15% for providers who joined before March 2016, and 20% for providers who joined after March 2016. The owner fee is zero if the owner joined before September 2015, while it varies but is never more than \$50 per booking for owners who joined after September 2015.¹⁴ DogVacay had a very similar fee structure and its commissions closely tracked those of Rover throughout the period between 2012 and 2017 (see Figure 3 in Section 4.2 below).

4.1 The Acquisition

On March 29, 2017, Rover announced it would buy DogVacay.¹⁵ Rover decided that it would shut down DogVacay and transfer all the business to the Rover platform rather than

¹²Details on how the current search algorithm works on Rover can be found at <https://www.rover.com/blog/sitter-resources/how-rover-search-works/> (accessed October 2020).

¹³At the time of our study, the only price suggestion available was Rover's "holiday rate" feature, which suggested sellers to increase their prices during holidays.

¹⁴Before July 2019, the maximum owner fee was \$25 per booking, according to screenshots on Wayback Machine. These screenshots can be accessed at <https://web.archive.org/web/20190705174452/https://support.rover.com/hc/en-us/articles/205385304-What-are-the-service-fees->. Information on current policies is available at <https://support.rover.com/hc/en-us/articles/205385304-What-are-the-service-fees-> (accessed December 2020).

¹⁵<https://techcrunch.com/2017/03/29/rover-dogvacay-merge/> (accessed July 2019).

maintaining both websites independently. DogVacay was reportedly struggling to keep up with the recent cash injections that Rover had received from venture capitalists,¹⁶ and Rover acquired DogVacay in an all-stock deal.¹⁷ Additional terms were not disclosed, but it is unlikely that the merger was subject to review by the Federal Trade Commission or the Department of Justice since the Hart-Scott-Rodino threshold for mandatory reporting was \$80.8 million in 2017. Neither the Federal Trade Commission nor the Department of Justice have a publicly available case involving Rover.¹⁸

Three features create a unique opportunity to study network effects from this acquisition: the acquisition led to a single aggregate platform; users migrated to the post-acquisition platform within 3 months; and we can identify the same users across the two platforms. First, it is rare for the acquired platform to merge with the acquiring platform. For example, even though Zillow acquired Trulia in 2015, the two platforms are still both active. The same is true for Google Maps and Waze, and for many online travel booking sites, such as Booking.com, Kayak, and Priceline, which are jointly owned by Booking Holdings. As Aaron Easterly, the CEO of Rover, confirms in a public interview,¹⁹ the decision to fully absorb DogVacay into the Rover brand was a consequence of the rapid growth that Rover was experiencing during the acquisition. At the time, Rover chose not to slow its growth to navigate the internal lobbying arising from two separate brands nor to integrate the back-ends while keeping two separate front-ends.

Second, the transfer of DogVacay's users to Rover happened quickly. In February 2017, Rover agreed to buy DogVacay. The acquisition was announced at the end of March. In early May, Rover announced that DogVacay would be shut down.²⁰ By early July, DogVacay ceased operations.

Third, when Rover announced that DogVacay would be shut down, Rover also started allowing DogVacay users to migrate their accounts to Rover. This meant that a user could

¹⁶<https://www.latimes.com/business/technology/la-fi-tn-dogvacay-rover-20170329-story.html> (accessed June 2020).

¹⁷<https://techcrunch.com/2017/03/29/rover-dogvacay-merge/> (accessed April 2020).

¹⁸<https://www.ftc.gov/news-events/media-resources/mergers-and-competition/merger-review> and <https://www.justice.gov/atr/merger-enforcement> (accessed April 2020).

¹⁹<https://soundcloud.com/acquiredfm/season-2-episode-10-the-rover> and <https://www.geekwire.com/2018/inside-rovers-dogvacay-deal-former-rivals-went-one-brand-not-two-acquisition/> (accessed April 2020).

²⁰Based on the publish date of this website: <https://www.rover.com/joining-forces/>

link their DogVacay account to their Rover account if they had been active on both platforms before the acquisition, or to a new Rover account otherwise. The account migration meant that a user would keep all their transaction and online rating history on the Rover platform. Among those users who did not actively migrate their accounts, multi-homing users could still be identified from their email address. While matching users on email addresses can sometimes be inaccurate, we are confident that the similarity of services exchanged on the two platforms likely incentivizes people who are serious about using both platforms to use the same email address.²¹

4.2 Data

We observe all service requests, buyer-seller booking inquiries, matches, and reviews from *both* platforms before and after the acquisition. A *request* refers to a buyer’s need for a sitter (e.g. dog boarding in Seattle from August 16th until August 18th) and is created when a buyer initiates a search or contacts a sitter directly. Contacts for the same request with different sellers are recorded as separate *booking inquiries*. A search leads to a recorded request only if a buyer sends at least one booking inquiry to a sitter. If a booking inquiry leads to a transaction, it is matched to a *stay*. Both DogVacay and Rover have multiple service categories, but we restrict attention to dog overnight boarding, which constitutes 70% of gross transaction volume on Rover and 91% on DogVacay before the acquisition.

We consider all buyer-seller booking inquiries initiated between June 2011 and January 2018 for requests between January 2012 and January 2018 included. Out of all booking inquiries, we remove those whose duration – i.e. number of nights requested – is recorded as negative or greater than 1 month (0.6% of requests), those with lead times – i.e. time between start date and request date – recorded as negative or greater than one year (1.1%), price outliers in terms of total price or commission fee percentage (2.3%). In particular, we remove prices lower than \$1 or higher than \$200 per night, and commission fees greater

²¹Survey evidence suggests that on average people have just less than two email accounts, and 2.5 when including a work account. Of those two accounts, one email address is often considered primary, and evidence suggests that there is huge inertia to changing the primary address. Finally, consumers are willing to share their primary address with businesses they trust. See <https://www.zettasphere.com/how-many-email-addresses-people-typically-use/>, which discusses results from the Data and Marketing Association (accessed April 2021).

than 30%. In total, we exclude 4.2% of total requests, and 3.8% of transactions.

The nature of competition between Rover and DogVacay before the acquisition suggests that a merger is likely to generate network effects. First, the two platforms were of similar size in the dog overnight boarding category before the acquisition, with Rover transacting at a 25% higher volume compared to DogVacay in the quarter before the acquisition.²² Second, the local nature of the services exchanged implies that buyers are typically interested in transacting with sellers within the same city. Indeed, 79% of booking inquiries, and 81% of stays occur within a buyer’s CBSA.²³ This means that we can measure competition between Rover and DogVacay at the local rather than aggregate level. Third, few users, and fewer buyers than sellers, multi-home across platforms. However, they account for a disproportionate share of transactions. We define a user as multi-homing if they transact at least once on both platforms over the 5 years before the acquisition. Only 3.3% of buyers and 7.6% of sellers multi-home. Not surprisingly, multi-homing users tend to transact more frequently than single-homing users. 27% of transactions are made by multi-homing sellers and 8% are made by multi-homing buyers.²⁴

Multi-homing sellers treat the two platforms as close substitutes, at least judging by the price they charge, even though DogVacay’s prices are higher on average.²⁵ Indeed, during the period before the acquisition, DogVacay sellers were expected to receive about \$3.50 more per night than sellers on Rover, or 13% more. After controlling for geographic and time observables, the price difference decreases to about 6% but it completely disappears once we compare prices of multi-homing sellers transacting on both Rover and DogVacay


²²Across all service categories, Rover was 62% larger than DogVacay. Appendix Figure D.2 plots the number of monthly stays on DogVacay since January 2012, in log scale. Despite being founded after Rover, DogVacay immediately outgrew Rover in overnight boarding services, before being surpassed again around March 2015.

²³CBSA stands for Core-Based Statistical Area, which roughly coincides with metropolitan and micropolitan areas.

²⁴Appendix Figure D.3 plots the share of a user’s transactions occurring on DogVacay prior to the acquisition, separately for buyers and sellers. On average, only 4.2% of users are both buyers and sellers of services on any given year. Buyers rarely act as service providers on the platforms. In the years before the acquisition, on average 4.8% of buyers also transacted as sellers on any given year. Sellers are more often buying pet-sitting services on the platforms. Indeed, 25.8% of sellers also transacted as buyers on any given year.

²⁵The payment that a seller receives is equal to what the buyer pays minus the platform commission fees. Tipping is not required, and is not recorded on the platform. However dog owners are not prevented from tipping sitters outside of the platform (<https://support.rover.com/hc/en-us/articles/206199686-Should-I-tip-my-sitter->, accessed July 2019).

Figure 3: Average Fees



The figure plots the average commission fee, as a percentage of the price that buyers pay. The vertical line identifies March 2017, when the acquisition was publicly announced. Levels on the y-axis are hidden to protect company information.

within the same month (Appendix Table D.1). This suggests that sellers on DogVacay may have different qualities or costs compared to sellers on Rover, which may induce sorting on the demand side as well. Nonetheless, multi-homing sellers consider buyers from the two platforms as close substitutes.

Figure 3 plots the average commission fee on the two platforms, computed as the ratio of platform total fees over the price paid by buyers. The figure shows that commission fees were very similar across platforms, and they continued their pre-acquisition upward trend after Rover acquired DogVacay. The upward trend is due to the higher fee schedule for buyers and sellers who joined after September 2015 and March 2016, respectively, whose shares increased steadily over time. As is clear from the figure, commission fees did not increase discontinuously after the acquisition, suggesting that Rover did not take advantage of its increased market power to capture a higher share of surplus.


5 Empirical Strategy and Identification

In this section, we describe how to test our theory, which relies on variation in pre-merger market shares across geographies. Figure 4 shows the distribution of Rover's market shares in 2016 (in terms of gross transaction volume) across zip codes with at least 50 stays in that year. Because buyers and sellers' zip codes may differ, we use sellers' zip codes for our

market definition. In the average zip code in 2016, Rover had about 53.6% market share, but there was substantial variation across zip codes. At least part of that variation can be explained by the different expansion strategies that Rover and DogVacay adopted years earlier when they just started out.²⁶

We separate zip codes into 5 groups: zip codes where in 2016 Rover had market shares below 20%; between 20% and 40%; between 40% and 60%; between 60% and 80%; and above 80%. Merging the two platforms after the acquisition was effective in migrating DogVacay users to Rover. Zip codes with Rover market shares smaller than 10% experienced a median increase in users on Rover of 550% while markets above 90% had a median increase of 14% (Appendix Figure D.4).

Figure 4: Rover Market Shares Pre-Acquisition



The figure plots the histogram of Rover market shares in 2016, the year prior to the acquisition. Each observation is a zip code with at least 50 transactions in 2016. The zip code’s Rover market share is defined using gross transaction volume.

Because the merger may have affected user outcomes beyond the direct effects of platform differentiation and network effects, and because our hypotheses from Section 3 vary by pre-merger market share, we can control for other changes affecting all geographies uniformly by using certain zip codes as a control group. The zip codes where Rover was already dominant pre-merger are the best candidates as control group, since Rover remained the

²⁶We find that part of the variation in 2016 market shares can be explained by which platform was the first mover in the market. Appendix Table D.2 shows that on average, Rover tends to have a 7% higher market share in zip codes where the first stay was booked on Rover rather than DogVacay. Due to confidentiality terms, we cannot disclose how the expansion strategies differed between Rover and DogVacay, although the two differed substantially in the way they targeted growth by expanding across geographies versus growing their user base within particular geographies.

dominant platform after the merger. We thus use zip codes where Rover had more than 80% of the pre-merger market as our control zip codes. Since Equation (1) is monotone in pre-merger market shares but Equation (4) is symmetric around 50% market share, we consider four separate treatment groups corresponding to the other four market share groups displayed in Figure 4.

Zip codes where either Rover or DogVacay were dominant before the acquisition tend to be more rural, have fewer residents, lower population densities, and lower shares of college graduates. Areas where Rover is particularly successful also tend to have higher pet ownership rates. Appendix Figures D.5 and D.6, together with Appendix Table D.3, provide comparisons for a large set of observable characteristics, platform performance metrics, and their evolution over time. Given these differences, we may be concerned that the main assumption behind a difference-in-differences approach, that zip codes with different market shares have the same latent trends in platform performance, does not hold.

To ensure that zip codes in treated market share groups are as similar as possible to zip codes in the control group, we employ a matching estimator that accounts for covariate imbalance across groups (Imai et al., 2018). We match one zip code from the control group to each “treated” zip code using covariate balancing propensity score matching (CBPS), introduced by Imai and Ratkovic (2014). Distances are calculated on the total number of active sellers in each month up to a year before the acquisition, where an active seller is defined as a seller who was involved in at least one booking inquiry in the given month. We hold the matched control group constant as we measure the effects of combining the two platforms across different outcomes of interest. Matching on number of sellers ensures that treated and control groups have similar number of participants across the two platforms combined.

Appendix Table D.4, which provides descriptive statistics for the matched samples, shows that we are able to improve matching on a number of covariates that we do not explicitly use in the matching procedure.²⁷ However, platform performance metrics that are not explicitly considered in matching (e.g. prices, match rates, and share of repeat transactions) fail to balance across treatment and control group. Some of this imbalance is

²⁷Appendix Table D.3 presents descriptives for the unmatched zip codes.

expected — for example we know that prices are higher on DogVacay and average prices will therefore be higher in markets with a higher DogVacay share. Other differences reflect the fact that platform performance metrics tend to positively correlate with a platform’s market share. We should note however, that our empirical strategy, described below, does not require identical levels of pre-treatment outcomes, but rather parallel trends. The figures in the results section below provide support for this assumption.

Given matched zip codes, let y_{zt} be the outcome in treated zip code z and year-month t . Separately for each treated market share group $[0 - 20\%)$, $[20\% - 40\%)$, $[40\% - 60\%)$, and $[60\% - 80\%)$, we estimate the following regression:

$$y_{zt} - y_{z't} = \alpha_t + \epsilon_{z,z',t}, \tag{5}$$

where z is the treated zip code, and z' is the matched control zip code. The coefficients α_t should be interpreted as changes in the outcome variable relative to the control group, and relative to February 2017, the month before the acquisition announcement. Cluster-robust standard errors are calculated using the method from Aronow et al. (2015).²⁸

Equation (5) allows us to test Equations (1) and (4). For Equation (1) network effects imply that the coefficients α_t after the merger should be positive and largest in market share group $[0 - 20\%)$. For Equation (4), if network effects are large enough to justify a single combined platform, we would expect the largest benefits from network effects to arise in the zip codes with shares $[40\% - 60\%)$.

To test Equation (3) we need a different approach. Recall that in order to evaluate the role of platform differentiation, we need to estimate to what extent DogVacay users are worse off *relative to* Rover users who experienced the same change in platform size. Rover users in markets with Rover’s pre-merger market share of s experience a change in platform size of $1 - s$. Symmetrically, DogVacay users in markets with Rover’s pre-merger market share of $1 - s$ experience an identical change in platform size of $1 - s$. We attribute any difference in outcomes between Rover and DogVacay users of these symmetric markets to

²⁸Each matched pair, or dyad, is no longer independently informative, as a single control market can impact the estimates of multiple dyads. The method proposed in Aronow et al. (2015) accounts for the correlation in error terms between each matched pair.

a reduction in platform differentiation.

Let $s \in \{0, 20\%, 40\%, 60\%, 80\%\}$ denote the lowest Rover’s market share in each of our market share groups. For each of the five s , we consider the outcomes of Rover users in zip codes with market shares within $[s, s + 20\%)$ and the outcomes of DogVacay users in zip codes with market shares within $[80\% - s, 100\% - s)$. With these outcomes we estimate the following regression:

$$y_{zt} = \beta_t + \gamma_t \mathbb{1}\{z \text{ has market share in } [80\% - s, 100\% - s)\} + \nu_z + \epsilon_{zt}, \quad (6)$$

where y_{zt} is the outcome of Rover users in zip code z and year-month t if $z \in [s, s + 20\%)$, or the outcome of DogVacay users in zip code z and year-month t if $z \in [80\% - s, 100\% - s)$.

The coefficients γ_t measure the difference in outcomes between DogVacay and Rover users in markets where both users experienced the same change in market size, and in month t relative to February 2017. Given Equation (3), we expect the γ to be negative due to the loss of platform differentiation.

In estimating Equations (5) and (6), we look at a large number of outcomes proxying for the various components of user surplus. Buyers’ utility is a function of the probability to find a sitter to transact with, the quality of the transaction, the price of the transaction, and search costs affecting the propensity of buyers to even post a request in the first place. We calculate the match rate of posted requests as the number of successful transactions in a given month and zip code divided by the number of posted requests. We compute the average nightly price of successful transactions as gross transaction volume divided by the number of transactions in a given month and zip code. We proxy for the average match quality with three metrics: the share of transactions in a given month and zip code whose buyer requests help again in the subsequent three months; the share of (non-repeat) transactions leading to a repeat stay in the future; and the share of transactions with a 5-star review submitted by the buyer. We also look at aggregate metrics: total number of unique buyers posting requests, total number of posted requests, and total number of transactions in a given month and zip code.

To test Equations (1) and (3), we need to categorize users as Rover users or DogVacay

users pre-merger. We define buyers as Rover buyers if all their booking inquiries during the year were on Rover, and DogVacay buyers if all their booking inquiries during the year were on DogVacay. We then measure the number of those users who post requests or engage in booking inquiries again in any given month of the following calendar year. The small share of multi-homers, those with inquiries on both platforms in a given year, are analyzed separately in Appendix B.

To test Equation (4) – are network effects large enough to justify a dominant platform? – we compute market-level outcomes by aggregating Rover and DogVacay outcomes (after DogVacay was shut down this will coincide with just Rover outcomes). We also ensure that the results hold true for new users, by focusing on outcomes for users who had never posted requests on any platform or market prior to the given month.

The combination of many outcomes and subsets of users for which to measure these outcomes (especially if we include analogous outcomes from the sellers’ perspective) would make our results section extremely long, but our main results can be summarized by just focusing on buyers and on two outcomes: request match rates and aggregate number of transactions. We present the other outcomes and a similar analysis for sellers in Appendix B. The results are also similar for more aggregated market definitions based on zip code clusters, which are less prone to potential violations of the stable unit treatment value assumption, but give rise to noisier estimates (Appendix B).

6 Results

This section presents our results.²⁹ We start with tests of platform level network effects (Equation 1). In this case, y_{zt} is the outcome of buyers in zip code z and year-month t for buyers who had posted booking inquiries only on Rover in the calendar year preceding t . Figure 5a plots the estimates of Equation (5) with log number of transactions and request match rates as the outcomes. As theory predicts, the top row shows that Rover buyers benefit more from the merger when the influx of users from DogVacay is larger. The effects

²⁹This section presents the results with event study plots. Appendix Tables B.1 through B.5 present the results of difference-in-differences regressions, aggregating the months in the pre-acquisition announcement period, those in between the announcement and the shut-down of DogVacay, and those after the shut-down of DogVacay.

on the top row imply a 26% increase in transactions for the markets with 0-20% market shares (first plot from the left) and around 17% increase in transactions for markets with 20-40% or 40-60% market shares (second and third plots). This increase in transactions is consistent with increased variety of sellers on the platform due to the migrating sitters from DogVacay. The increase in activity from Rover buyers entirely comes from the extensive margins – more users posting requests – rather than match quality or match rates. Indeed the bottom row of Figure 5a shows that Rover buyers did not experience an improvement in match rates, and the appendix confirms that our proxies for match quality remain unchanged (Appendix Figure B.1).

The top row of Figure 5a is our evidence that network effects exist and provides a potential justification for merging two platforms. Buyers of the acquiring platform benefit from the influx of users from the acquired platform as evidenced by their increased propensity to post requests (and hence engage in more transactions) and by the fact that this increased propensity to post requests is larger in markets with a bigger shock to the number of users.

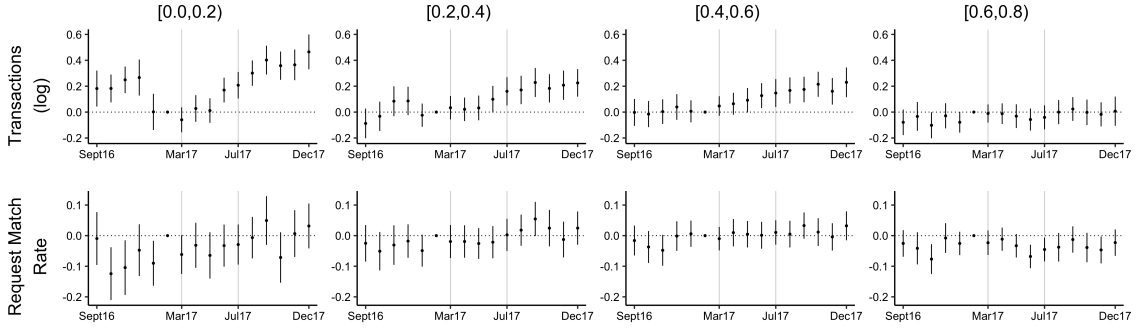
Next, we evaluate the effect of the merger on DogVacay users. We start with Figure 5b, which mirrors the previous analysis. Here, y_{zt} is the outcome of buyers in zip code z and year-month t for buyers who had posted booking inquiries only on DogVacay in the calendar year preceding t . The top row of Figure 5b shows that DogVacay buyers experience higher attrition and lower match rates compared to before the acquisition and compared to DogVacay buyers in the zip-codes where Rover had 80-100% market share prior to the acquisition.³⁰

The negative coefficients in Figure 5b seem to suggest that DogVacay buyers are worse off after the merger. However, the $v(1) - v(1 - s)$ component in Equation (2) is positive and largest in the control markets, which implies that DogVacay buyers in control markets experience the largest network effect benefits from combining with Rover users. This explains why DogVacay buyers in the treatment markets are worse off compared to DogVacay buyers in the matched control markets.


To evaluate the role of platform differentiation, we need to estimate Equation (6). The γ coefficients in that equation represent to what extent DogVacay buyers are worse off *relative*

³⁰ Appendix Figure B.2 shows that the reduction in transactions is largely due to a reduction in the number of buyers rather than the frequency of transactions per transacting buyer.

Figure 5: Estimates of Merger Effects at the Platform Level



(a) Rover Users

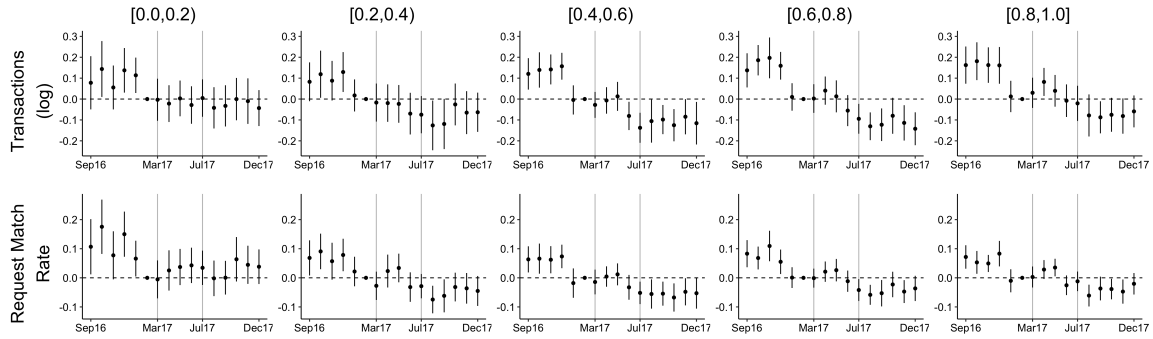


(b) DogVacay Users

Regression estimates of Equation (5). In the first panel we test Equation (1). The first row displays results where the outcome is the (log) number of transactions from buyers who, in the prior calendar year, had only engaged in booking inquiries on Rover. The second row displays results for the match rate of those same Rover buyers, *i.e.*, the number of stays divided by the number of requests posted by existing Rover buyers. Panel (b) displays analogous outcomes for users who, in the prior year, had only engaged in booking inquiries on DogVacay. An observation is a matched zip code-month. In each panel the regressions come from 2 different outcomes — stays and match rates — and 4 treatment groups — zip codes with Rover’s market shares in the following bins: 0-20%, 20%-40%, 40%-60%, and 60%-80%. The control group from which matched zip codes are selected includes zip codes with Rover’s market shares greater than 80%. Grey vertical lines denote March and July 2017, the months when the acquisition was announced and DogVacay was effectively shut down, respectively. Extensions, including other outcomes, results for multi-homing users, and estimates with clusters of zip codes as markets are in Appendix B.

to Rover buyers who experienced the same potential change in platform size. Figure 6 plots the estimated γ coefficients for each month leading up to and after the acquisition. Across all market share groups, it is clear that DogVacay buyers experienced a reduction in the number of transactions (top row) and request match rate (bottom row) relative to Rover buyers in symmetric markets around the merger announcement. In fact, the decline in outcomes started occurring in January and February 2017, before the merger was announced

Figure 6: Estimates of Merger Effects For DogVacay Users Relative to Rover Users



Regression estimates of Equation (6) testing Equation (3). The first row displays results where the outcome is the (log) number of transactions from buyers who, in the prior calendar year, had only engaged in booking inquiries on Rover or DogVacay. The second row displays results for the match rate of those same users, i.e., the number of stays divided by the number of requests submitted. Each column corresponds to a market share group ($s, s + 20\%$). Given ($s, s + 20\%$) the figure plots the estimated difference in outcomes between DogVacay users in markets with Rover market shares in $(80\% - s, 100\% - s)$ and Rover users in markets with Rover market shares in $(s, s + 20\%)$. So for example, the top-left plot compares the (log) number of transactions that DogVacay users exchanged in markets where Rover had pre-merger market shares above 80% and the number of transactions that Rover users exchanged in markets Rover had pre-merger market shares below 20%.

but presumably during merger talks. This decline continued during the March-July 2017 period as DogVacay users started migrating to Rover. Outcomes drop more drastically after DogVacay was shut down and then stabilize. After DogVacay was shut down, the reduction in transactions of DogVacay buyers relative to Rover is at least 10% across all market share groups, and the reduction in match rates is at least 4 percentage points (Appendix Table B.3).

Equation (3) predicts that this decline is decreasing in absolute value as Rover's market share increases. Our empirical results are not fully aligned with that prediction. In fact, the largest declines in transactions and match rates are experienced in the markets with Rover shares between 20% and 80%. A model extension, which allows for a share of DogVacay users to choose the outside option instead of migrating to Rover and which is consistent with our empirical findings, can explain this pattern (see Appendix A).

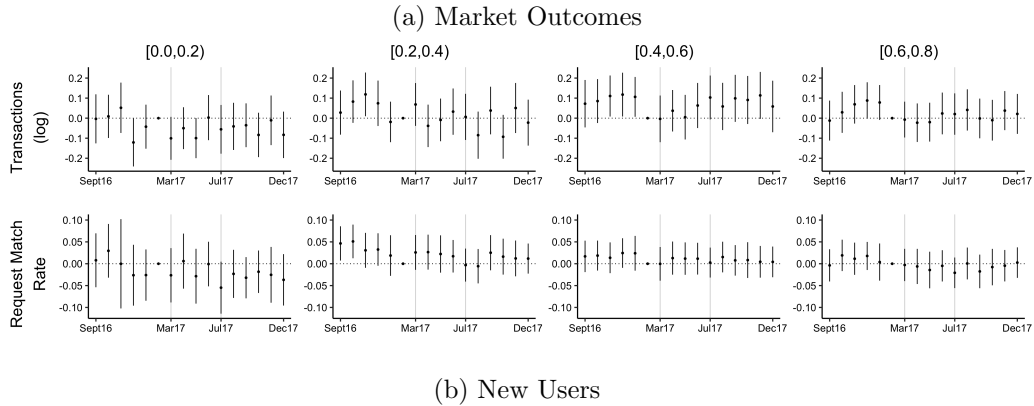
The final step in our analysis is to test whether network effects are large enough that they more than offset the decline in platform differentiation. Our theory provides Equation (4) as a test. Figure 7a plots the results of the test. The outcome in the first row is the (log)

total number of transactions in a given zip code-month, regardless of whether they were intermediated by DogVacay or Rover. Like before, each column corresponds to a different treatment group. This time however, if network effects dominate the reduction in platform differentiation, we would expect the largest increase in the number of transactions to occur in the zip codes where Rover’s market share was between 40% and 60% pre-merger (third plot in the first row). The effect should then be monotonically decreasing for the plots to the right and to the left. With symmetry of the two platforms, we expect that the group with Rover’s market share between 0 and 20% (first plot from the left) to be indistinguishable from the control group. These patterns should be true not only for transactions, but also for request match rates (bottom row).

The first row of Figure 7a shows that indeed, there seems to be an uptick in the number of transactions after merging the two platforms in the zip codes with 40-60% market shares, but the estimated effect is noisy and often is indistinguishable from a null effect. Pooling together the months after DogVacay’s shutdown to estimate a single difference-in-differences coefficient for each treatment group (Appendix Table B.4) confirms that the effect is not statistically significant. Zip codes with market shares farther away from 40%-60% are indistinguishable from the control group and, if anything, the difference-in-differences coefficient for 0-20% and 20-40% market share groups implies a marginally significant 7.5% decrease in the number of transactions. Similarly for the request match rate (second row of Figure 7a), we don’t find any positive effect of the merger across market share groups. For zip codes where Rover had less than 20% market share, we even find a significant reduction in match rates of 3.5 percentage points. These results and the results in Appendix Figure B.4 suggest that buyers do not find matches of higher quality or at higher rates with the single merged platform compared to when there were two competing platforms.

The same conclusion rejecting the hypothesis that a single platform is better for users than two competitors is true when focusing on new users only. Figure 7b displays regression estimates of Equation (5) using transactions and match rates of new buyers, defined as those who never posted a request or were involved in a booking inquiry on any platform prior to the current month. The plots show surprisingly stable transaction volumes and match rates after the merger across all treatment groups relative to the control group. This is a

Figure 7: Net Effects at the Market Level



Regression estimates of Equation (5) to test Equation (4). Panel (a) presents market-level outcomes (log transactions and request match rate), while Panel (b) focuses on the same outcomes for new users, defined as users who never had a booking inquiry before the given month. Otherwise the plots are identical to Figure 5. Extensions and robustness checks are in Appendix B.

notable result, because it shows that horizontal preferences for platforms are not something that users develop *after* joining a particular platform, which would lead us to finding that a single dominant platform is on average preferred by new users than two competitors.

6.1 Overview of Robustness Checks

In recent years there have been concerns that the specific implementation of a difference-in-differences strategy may affect the outcome. For example, Orchinik and Remer (2020) show that synthetic controls, two-way fixed effects, and matching sometimes give differing results for the effect of airline mergers. Our preferred specification throughout the paper uses a combination of matching and difference-in-differences. We have also conducted our analysis using several other strategies, which we briefly describe below.

The first alternative strategy we implement is the synthetic difference-in-difference estimator of Arkhangelsky et al. (2021). This estimator combines several desirable features of synthetic control estimators and difference-in-differences estimators. Like the synthetic control method, it creates a comparison group that matches pre-treatment trends between treatment and control groups. Like difference-in-differences, it allows for unit-level shifts (i.e. differences in levels between treatment and control groups) and inference with many treated units. Lastly, the estimator flexibly estimates weights placed on pre-treatment time-periods, reducing concerns regarding researcher flexibility in the choice of pre-treatment period length.

We describe the details of our implementation of the synthetic difference-in-differences estimator in Appendix B and the results are shown in Figure B.7. These results are largely consistent with our main findings. Outcomes substantially improve for Rover users but not as much at the market level. At the market level, we find a small increase in transactions for the 40% to 60% group but find no improvements in the match rate.

As an alternative robustness check, we estimate a difference-in-differences specification without matching. We once again find similar results. More details about the unmatched specification are available in Appendix B and the results are shown in Figure B.8.

6.2 Heterogeneous Effects Across Markets and Users

The results so far confirm the existence of network effects at the platform level but that the benefits from network effects are not enough for the average consumer to prefer a single combined platform over two separate platforms. In fact, our results suggest that the two options are, at least in the short-run, similar for consumer surplus. This is true on average, but network effects may dominate over platform differentiation in certain markets and not in others. We explore two dimensions of heterogeneity across markets: market size, and propensity to multi-home. We leave the theoretical discussion to Appendix A and histograms of these characteristics across zip codes in Appendix Figure D.7.

Markets differ in their total number of transactions. Among zip codes with at least 50 transactions in 2016, the average zip code had 171 stays in the same year, but with a standard deviation of 146, demonstrating that there is substantial heterogeneity across zip

codes. It is possible that both platforms were already operating at an efficient scale in large markets but not in small markets. If this were the case, we would expect platform-level network effects to be larger after the merger in smaller markets. It would also be more likely that network effects are large enough to justify a single platform in smaller markets compared to larger markets.

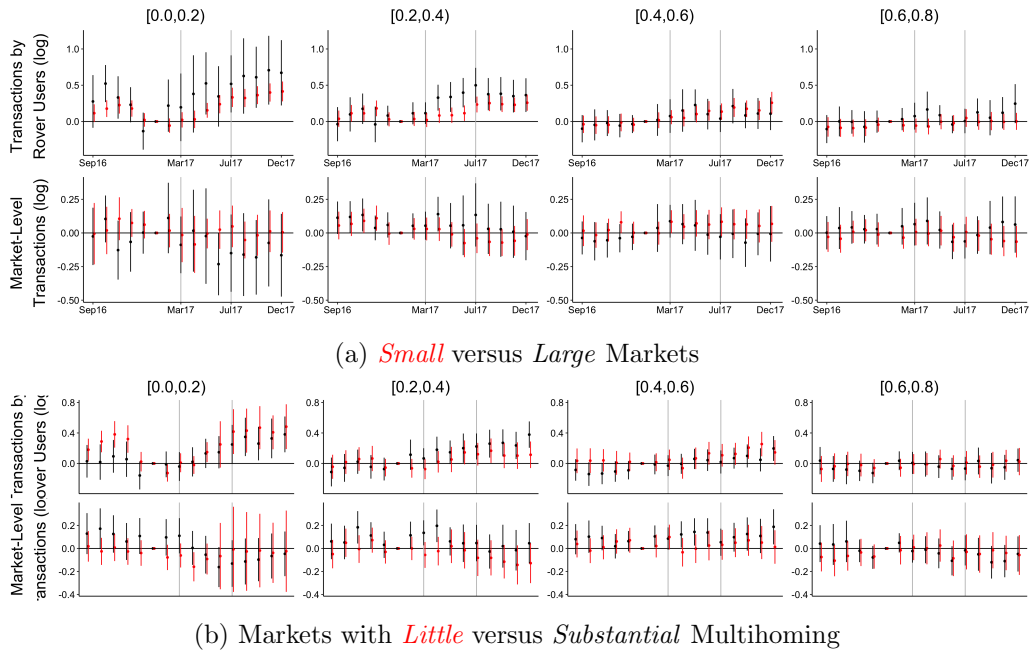
To test these hypotheses, we split zip codes into those with more or fewer than 250 transactions in 2016. A large share, 81%, of our zip codes are considered small markets. Figure 8a plots the estimates testing Equation (1) on the top row and Equation (4) on the bottom row.³¹ The red estimates are for small markets, while the black estimates are for large markets. We do not find much of a statistically significant difference between small and large markets: the transactions by existing Rover users go up monotonically in the influx of new users from the acquired platform (top row), while the market-level transactions do not increase relative to the control group (bottom row).

We should note that there is likely to be selection into the markets that experienced a small versus large number of transactions in 2016, so we have only suggestive evidence that our results are independent of market size. In other words, it may very well be that markets that were large in 2016 would have benefited more from a merger had it happened earlier, when the two platforms were smaller.

A second dimension of heterogeneity is the propensity to multi-home. In the extreme case and without capacity constraints, if one side of users – buyers or sellers – fully multi-home, every user has access to every other user in the market, so combining the two platforms should have no effect on the number and type of exchanges available to each user. Given that sellers are more likely to multi-home, we look at differential effects of merging the two platforms by sellers' propensity to multi-home. In the average zip code, 23.5% of 2016 transactions were supplied by multi-homing sellers, but the standard deviation of 29.9% suggests there is substantial variation. We separate zip codes at the 10% cutoff, i.e., where 10% of transactions are supplied by multi-homing sellers. About half of the zip codes are at each side of the cutoff. We would expect larger benefits from merging the two platforms to occur in markets with a smaller propensity to multi-home.

³¹Note that the matched samples differ from Figures 5 and 7 because we constrain each treated zip code to be matched to a control zip code within the same market size group.

Figure 8: Estimates of Merger Effects – Heterogeneity by Market Type



Estimates of Equation (5) for different markets. In Panel (a) the zip codes are divided into two groups: markets with 50-250 transactions in 2016 (in red), and markets with more than 250 transactions (in black). In Panel (b), the zip codes are divided into two groups: markets whose share of 2016 transactions completed by multi-homing sellers is less than 10% (in red) and those whose share of transactions by multi-homing sellers is greater than the cutoff (in black). We focus on log number of transactions by buyers who only used Rover in the preceding calendar year (top row in each panel) and market-level log number of transactions (bottom row in each panel). Across all panels, coefficients in red denote zip codes where we would expect the improvements from the merger to be bigger (as per our theory extensions in Appendix A). Otherwise each panel is identical to the top rows of Figure 5a and Figure 7a.

Figure 8b displays the results of matched sample regressions for markets with low propensity (red) and high propensity to multi-home (black). Similarly to what we found for market size, transactions by Rover users increase monotonically in the market share of the acquired platform, while market-level transactions are similar between treatment and control groups, regardless of sellers' propensity to multi-home.

6.3 Predictors of User Behavior Post-Merger

In this subsection, we offer some suggestive evidence on potential reasons why DogVacay users, despite benefiting from the increase in platform size, are worse off relative to Rover users. To do so, we match DogVacay users to similar Rover users and explore their activity after DogVacay was shut down. We find that DogVacay buyers have 30% fewer transactions relative to similar Rover buyers after DogVacay was shut down. Although switching costs partially explain this difference, we find support for two alternative explanations. The first is that DogVacay buyers may continue transacting with prior partners off platform (disintermediation). The second is that there may be coordination failures, so that DogVacay buyers are not able to find prior transaction partners on Rover, in part because those sellers do not switch platforms.

Our empirical strategy is based on matching Rover and DogVacay buyers who had at least one transaction in 2016, based on their activity throughout 2016. Each buyer is associated to a unique market corresponding to the modal zip code of the sellers with whom they communicated in 2016. We use coarsened exact matching on the number of transactions and booking inquiries, the month of the last transaction, whether they had at least one transaction in 2016 with a repeat seller, the average nightly price of all their 2016 transactions, and Rover's pre-merger market share in their market. For the latter, we match DogVacay buyers from market share group $[80\% - s, 100\% - s)$ with Rover users from market share group $[s, s + 20\%)$. We then conduct regression analysis using the matching weights and excluding users for whom there was no match (Hong (2010)). Our outcome of interest is a buyer's total number of transactions between August and December 2017, after DogVacay was shut down.

Our first result, displayed in Table 1, column (1), shows that even among users with

Table 1: Transactions of Buyers After DogVacay is Shut Down

	# Transactions Post DogVacay Shutdown			
	(1)	(2)	(3)	(4)
DogVacay User	-0.2234*** (0.0065)	-0.0978*** (0.0057)	-0.1504*** (0.0101)	-0.0346** (0.0152)
# 2016 Stays	0.0750*** (0.0033)	0.0802*** (0.0044)	0.0804*** (0.0044)	0.1370*** (0.0083)
Avg. Nightly Price (2016)	0.0016*** (0.0002)	0.0016*** (0.0002)	0.0016*** (0.0002)	0.0016*** (0.0002)
Has Repeat Stay		0.0727*** (0.0129)	0.0729*** (0.0129)	-0.0846*** (0.0199)
DogVacay User × Has Repeat Stay		-0.2381*** (0.0126)	-0.3899*** (0.0204)	-0.1071*** (0.0286)
DogVacay Seller Migrated			0.0634*** (0.0103)	0.0622*** (0.0102)
Has Repeat Stay × DogVacay Seller Migrated			0.1712*** (0.0200)	0.1716*** (0.0185)
DogVacay User × # 2016 Stays				-0.0937*** (0.0093)
Mean of Y	0.74	0.74	0.74	0.74
R ²	0.02732	0.02928	0.03022	0.03509
Observations	212,817	212,817	212,817	212,817
Month of Last Stay FE	✓	✓	✓	✓
Platform Share FE	✓	✓	✓	✓

This table displays coefficients of regressions where the outcome is the number of transactions of a user post-DogVacay shut-down. Each observation is a single-homing buyer who had at least one transaction in 2016. The control variables include whether the user was on DogVacay in 2016, the number of stays in 2016, the average nightly price, whether a stay in 2016 was a repeat stay with a sitter from a prior transaction, and whether the seller migrated their profile to Rover post-merger (only applies to DogVacay users). A similar analysis for sellers is presented in D.5.

similar pre-merger behavior, DogVacay users are less likely to transact after DogVacay is shut down. The average number of transactions is 0.74, and DogVacay buyers engage in 0.22 fewer transactions relative to Rover buyers. This effect is economically important, representing an almost 30% drop in transactions. The next columns in Table 1 break down the drop across a few potential explanations.

The first potential explanation is that dog owners prefer to engage in repeat transactions with prior sellers. This is a characteristic of many platforms for local services, like food delivery or childcare, although it does not apply to all platforms, such as ride-sharing. On average, 50.8% of 2016 transactions are between a buyer and a seller who had already transacted with each other before. If buyers and sellers trust each other, then they may be willing to disintermediate the platform and avoid paying commission fees. The shutdown

of DogVacay could thus lead some users to disintermediate rather than migrating to Rover.

If disintermediation occurs, then we would expect that DogVacay buyers with repeat transactions would have fewer post-shutdown transactions relative to similar Rover buyers. Indeed, column (2) of Table 1 finds that this is the case. DogVacay buyers with a prior repeat stay have 0.24 fewer transactions post-shutdown compared to Rover buyers with prior repeat stays. For Rover buyers instead, having a prior repeat stay is positively correlated with subsequent transactions. Taken together, the estimates imply that while Rover buyers with repeat stays engage in more transactions in the second half of 2016 relative to buyers without repeat stays – consistent with repeat stays being a positive quality signal that the platform is providing value to its users – DogVacay buyers with repeat stays actually engage in fewer transactions than buyers without repeat stays, perhaps due to disintermediation.

Another explanation for why DogVacay users are worse off relative to Rover users is that DogVacay users may not be able to find each other on Rover. Both buyers and sellers need to migrate to the acquiring platform, but not all DogVacay sellers migrated to Rover. Buyers who did not find their prior sitter may have been induced to stop searching or send a request to a less preferred sitter. If this were true, then sellers' decisions to join Rover would help predict the transactions post-merger of the buyers with whom they interacted before the merger.

To study this coordination failure, we measure whether a DogVacay buyer's last seller in 2016 migrated their account on Rover post-merger. We add this dummy variable in column (3) of Table 1. We see that DogVacay buyers have 0.06 more transactions on Rover if their most recent DogVacay seller migrated, and that having a prior repeat stay and a seller who migrated is associated with an additional 0.17 increase in the number of transactions. This is consistent with the presence of coordination failures.

Finally, we show that frequent users of DogVacay, who presumably have a high value from dog-sitting are also hurt by DogVacay's shutdown. In column (4) of Table 1, we add an additional predictor: we interact an indicator for whether a buyer was on DogVacay in 2016 with their number of transactions in 2016. We find a negative coefficient of 0.09, implying that the more active DogVacay buyers had fewer transactions after DogVacay was shut down relative to similar Rover users. This negative coefficient is not as large as the

0.14 coefficient on the number of stays. This suggests that switching costs at least partially explain DogVacay users' attrition since high value DogVacay users have the most incentives to switch platforms

To summarize, DogVacay buyers use the combined platform less than similar Rover buyers after DogVacay was shut down. This drop seems to be at least in part explained by the role that repeat transactions play in the market for pet-sitting services, the ease with which buyers and sellers can transact off the platform, and switching costs.

7 Conclusions

There is a heated debate over the best way to operate digital platforms and how these platforms should be regulated (Scott Morton et al., 2019). A key tradeoff in this debate is whether multiple platforms are superior to single platforms, in terms of the value generated to the users. On one hand, competition gives users more options, which can benefit markets where users have differentiated preferences over platform features or the types of users they'd like to interact with. Competition can also result in lower prices and more innovation. On the other hand, if network effects are large, it may be best to have all users participating on a single platform.

In this paper, we show that while the risk of higher prices may not materialize in the short-term (perhaps because these companies were still growing quickly), platform differentiation can be an important factor offsetting network effects even in industries where competing platforms appear very close substitutes. Using the merger of the two largest platforms for pet-sitting services into a single platform, we evaluate how combining two platforms differentially affects markets that were already effectively experiencing a single platform—because the acquiring platform already had over 80% of the market—versus markets where the two platforms were competing on equal grounds or where the acquired platform was dominant.

We find that the acquiring platform experienced sizable network effects. Indeed, existing users of the acquiring platform increased their platform activity, more so in locations experiencing a bigger influx of users from the acquired platform. Although network ef-

fects are often assumed to exist in digital platforms, we provide one of the few empirical confirmations of their existence.

Despite the network benefits at the platform level, we find that on average at the market level, users are equally well off with one or two platforms, as evidenced by the constant number of transactions, match rates, and proxies for match quality. Combined with our evidence that platform prices did not increase post-acquisition, our results suggest that, on average, a single platform does not provide larger consumer surplus than the sum of two competing platforms.

We find that while existing users of both platforms benefit from aggregating user participation on a single platform, users of the acquired platform are worse off relative to users of the acquiring platform. In particular, they match at lower rates and complete fewer transactions. We show that some of this difference is likely driven by the importance of repeat transactions (which may lead to disintermediation) and switching costs. However, our finding that even for new users a single platform is not better than two competitors suggests that horizontal preferences do not simply originate from experience in using a particular platform.

Our study focuses on two platforms that intermediate local and time-sensitive services. Other platforms with similar features include ride-sharing (Lyft), food delivery (Doordash), home-improvement (HomeAdvisor), and child care (Care.com). These platforms are well suited for a similar causal analysis of network effects because they are comprised of geographically separate markets exchanging services. Our analysis of user attrition post-merger shows that repeat transactions play an important role counterbalancing network effects. As a result, platform differentiation may be even more important on child care platforms than in our setting, where repeat transactions are more frequent, and less important on ride-sharing platforms, where repeat transactions are rarer.

These results have important implications for both platform strategy and for antitrust regulation. From a strategic perspective, we show that it may be beneficial for a company to operate multiple platforms rather than combine them, offering a novel rationale for the many instances of platform acquisitions that kept acquired platforms in operation (e.g., Zillow and Trulia, or the many online travel sites within the Booking Holdings group). Owners

of multiple platforms have to balance network effects and scale economies of operating a single platform versus users' preferences for differentiation. When platform differentiation dominates, platform owners can harness network effects by making it easier for users to multi-home, for example by automatically cross-listing explicitly consenting users across platforms.

From an antitrust perspective, user outcomes were comparable between two platforms and a single platform in our setting. Together with the fact that platform commission fees did not increase after the acquisition and that kennels and dog hotels still constitute a large share of the market for pet-sitting services, our results point to the merged platform being better able to compete with larger incumbents by reducing fixed and variable costs, such as technology investments, customer acquisition costs, and labor costs. These considerations would of course be different in a context where the acquiring platform were the only option to access pet-sitting services.

The null effect at the market level occurs despite the presence of network effects that exist at the platform level, and despite the fact that the two platforms appear similar in the way they intermediate services. In other contexts where mergers occur between platforms that are not as close substitutes, horizontal preferences and user attrition are likely to play an even bigger role when comparing a single dominant platform versus multiple competitors. In those cases, it may be particularly important to ensure platform competition.

The strategic and antitrust angles are interrelated. Our study shows that network effects may not lead to winner-take-all equilibria – after all, DogVacay was acquired while it was growing similarly to Rover. Network effects are also unlikely to serve as justification for monopoly platforms operating in local services. Strategic competitors' acquisitions thus have to be motivated by other benefits in addition to network effects (such as scale economies) or may be allowed only with restrictions on the degree of company integration.

Given the difficulty of causally linking a merger in 2017 to events occurring many months later, our evidence concerns just the short-run. Extending theories and empirical approaches to estimate the effects of mergers beyond the first few months would be crucial to assess the costs and benefits of acquisitions of early stage competitors by incumbent platforms.

We have also focused on local, as opposed to global effects. Many important platforms

also enjoy global network effects across geographies, such as platforms for virtual work like Upwork, or app platforms like iOS and Android. Our paper does not speak to whether it is better for consumers to have two platforms with non-overlapping geographic presence or a single platform active in all geographies (Zhu et al., 2019), nor are we able to measure cost efficiencies from the acquisition. We leave the exploration of these topics for future research.

References

- ARKHANGELSKY, D., S. ATHEY, D. A. HIRSHBERG, G. W. IMBENS, AND S. WAGER (2021): “Synthetic Difference-in-Differences,” American Economic Review, 111, 4088–4118. 30, 49
- ARMSTRONG, M. (2006): “Competition in two-sided markets,” The RAND Journal of Economics, 37, 668–691. 6
- ARONOW, P. M., C. SAMII, AND V. A. ASSENOVA (2015): “Cluster-robust variance estimation for dyadic data,” Political Analysis, 564–577. 22
- BAKOS, Y. AND H. HALABURDA (2019): “Platform competition with multihoming on both sides: Subsidize or not?” Working Paper. 7, 8
- BERRY, S. T. AND J. WALDFOGEL (1999): “Free entry and social inefficiency in radio broadcasting,” RAND Journal of Economics, 30, 397–420. 8
- BOUDREAU, K. (2010): “Open platform strategies and innovation: Granting access vs. devolving control,” Management science, 56, 1849–1872. 7
- CAILLAUD, B. AND B. JULLIEN (2003): “Chicken & egg: Competition among intermedia-tion service providers,” RAND journal of Economics, 309–328. 6, 7
- CHANDRA, A. AND A. COLLARD-WEXLER (2009): “Mergers in two-sided markets: An application to the Canadian newspaper industry,” Journal of Economics & Management Strategy, 18, 1045–1070. 8
- CHAVENT, M., V. KUENTZ-SIMONET, A. LABENNE, AND J. SARACCO (2018): “Clust-Geo: an R package for hierarchical clustering with spatial constraints,” Computational Statistics, 33, 1799–1822. 50, 51
- CHU, J. AND P. MANCHANDA (2016): “Quantifying cross and direct network effects in online consumer-to-consumer platforms,” Marketing Science, 35, 870–893. 8
- CULLEN, Z. AND C. FARRONATO (2021): “Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms,” Management Science, 67, 3985–4003. 3, 8

- CUSUMANO, M. A., A. GAWER, AND D. B. YOFFIE (2019): The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, and Power, HarperBusiness. 2
- DUBÉ, J.-P. H., G. J. HITSCH, AND P. K. CHINTAGUNTA (2010): “Tipping and concentration in markets with indirect network effects,” Marketing Science, 29, 216–249. 9
- FARRELL, J. AND G. SALONER (1985): “Standardization, compatibility, and innovation,” Rand Journal of Economics, 16. 6
- FARRONATO, C. AND A. FRADKIN (Forthcoming): “The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry,” . 8
- FONG, J. (2019): “Search, Selectivity, and Market Thickness in Two-Sided Markets: : Evidence from Online Dating,” Working Paper. 8
- FRADKIN, A. (2018): “Search, matching, and the role of digital marketplace design in enabling trade: evidence from Airbnb,” Working Paper. 8
- GANDAL, N. (1994): “Hedonic price indexes for spreadsheets and an empirical test for network externalities,” the RAND Journal of Economics, 160–170. 7
- GOWRISANKARAN, G. AND J. STAVINS (2004): “Network externalities and technology adoption: lessons from electronic payments,” RAND Journal of Economics, 35, 260–276. 8
- GREENSTEIN, S. M. (1993): “Did installed base give an incumbent any (measureable) advantages in federal computer procurement?” The RAND Journal of Economics, 19–39. 7
- HAGIU, A. AND J. WRIGHT (2014): “Marketplace or reseller?” Management Science, 61, 184–203. 7
- HALABURDA, H., M. JAN PISKORSKI, AND P. YILDIRIM (2018): “Competing by restricting choice: The case of matching platforms,” Management Science, 64, 3574–3594. 15

- HE, L., I. REIMERS, AND B. SHILLER (2021): “Does Amazon Exercise its Market Power? Evidence from Toys R Us,” Working Paper. 9
- HOLTZ, D. M. AND S. ARAL (2018): “Limiting Bias from Test-Control Interference in Online Marketplace Experiments,” Working Paper. 50
- HONG, G. (2010): “Marginal mean weighting through stratification: Adjustment for selection bias in multilevel data,” Journal of Educational and Behavioral Statistics, 35, 499–531. 33
- IMAI, K., I. S. KIM, AND E. WANG (2018): “Matching methods for causal inference with time-series cross-section data,” Working Paper. 21
- IMAI, K. AND M. RATKOVIC (2014): “Covariate balancing propensity score,” Journal of the Royal Statistical Society: Series B (Statistical Methodology), 76, 243–263. 21
- JEZIORSKI, P. (2014a): “Effects of Mergers in Two-Sided Markets: The US Radio Industry,” 6, 35–73. 8
- (2014b): “Estimation of Cost Efficiencies from Mergers: Application to US Radio,” 45, 816–846. 8
- KABRA, A., E. BELAVINA, AND K. GIROTRA (2017): “Designing promotions to scale marketplaces,” Working Paper. 8
- KATZ, M. L., C. SHAPIRO, ET AL. (1985): “Network externalities, competition, and compatibility,” American economic review, 75, 424–440. 6
- KAWAGUCHI, K., T. KURODA, AND S. SATO (2020): “Merger Analysis in the App Economy: An Empirical Model of Ad-Sponsored Media,” . 9
- LAM, C. T., M. LIU, AND X. HUI (2021): “The geography of ridesharing: A case study on New York City,” Information Economics and Policy, 100941. 8
- LI, H., G. ZHAO, R. JOHARI, AND G. Y. WEINTRAUB (2021): “Interference, Bias, and Variance in Two-Sided Marketplace Experimentation: Guidance for Platforms,” arXiv preprint arXiv:2104.12222. 50

- LI, J. AND S. NETESSINE (2020): “Higher market thickness reduces matching rate in online platforms: Evidence from a quasiexperiment,” Management Science, 66, 271–289. 8
- NIKZAD, A. (2020): “Thickness and Competition in On-demand Service Platforms,” Working Paper. 7
- ORCHINIK, R. AND M. REMER (2020): “What’s the Difference? Measuring the Effect of Mergers in the Airline Industry,” Measuring the Effect of Mergers in the Airline Industry (May 15, 2020). 29
- PARKER, G. G. AND M. W. VAN ALSTYNE (2005): “Two-sided network effects: A theory of information product design,” Management science, 51, 1494–1504. 6
- RESHEF, O. (2019): “Smaller Slices of a Growing Pie: The Effects of Entry in Platform Markets,” . 8
- ROCHET, J.-C. AND J. TIROLE (2003): “Platform competition in two-sided markets,” Journal of the european economic association, 1, 990–1029. 2, 6
- RYSMAN, M. (2004): “Competition between networks: A study of the market for yellow pages,” The Review of Economic Studies, 71, 483–512. 7, 8
- SALONER, G. AND A. SHEPARD (1995): “Adoption of technologies with network effects: an empirical examination of the adoption of automated teller machines,” The RAND Journal of Economics, 26, 479–501. 7
- SCOTT MORTON, F., P. BOUVIER, A. EZRACHI, B. JULLIEN, R. KATZ, G. KIMMELMAN, A. MELAMED, AND J. MORGENSTERN (2019): “Committee for the Study of Digital Platforms: Market Structure and Antitrust Subcommittee-Report,” Chicago: Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business. 36
- SEAMANS, R. AND F. ZHU (2014): “Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers,” Management Science, 60, 476–493. 8
- TAN, G. AND J. ZHOU (2020): “The Effects of Competition and Entry in Multi-sided Markets,” The Review of Economic Studies. 7

TUCKER, C. (2008): “Identifying Formal and Informal Influence in Technology Adoption with Network Externalities,” Management Science, 54, 2024–2038. 8

WEYL, E. G. (2010): “A price theory of multi-sided platforms,” American Economic Review, 100, 1642–72. 6, 7

ZHU, F. AND M. IANSITI (2012): “Entry into platform-based markets,” Strategic Management Journal, 33, 88–106. 7

ZHU, F., X. LI, E. VALAVI, AND M. IANSITI (2019): “Network Interconnectivity and Entry into Platform Markets,” Harvard Business School Working Paper, 19–06. 39