Bouncing with the Joneses? A neural network approach to consumption neighborhood effects

Erik Grenestam*

January 29, 2019

EARLY DRAFT VERSION

Abstract

If one of your neighbors decides to purchase a good, does that affect your decision to buy the same good? I study this question using household level data collected in a novel way. I use an image classification algorithm to process a large set of aerial photos in order to infer household ownership of a visible durable good, specifically a trampoline. To estimate the neighborhood effect, I use a neighborhood fixed effects approach together with exogenous variation stemming from new neighbors moving in. I find that neighborhood effects are present, but only within short distances. The effect is relatively modest in size and has a marginal impact on aggregate demand for this particular product.

*Department of Economics and Center for Economic Demography, Lund University School of Economics and Management. I am grateful for financial support from Jan Wallanders and Tom Hedelius foundation and Stiftelsen för främjande av ekonomisk forskning vid Lunds universitet. I would also like to thank Mats Högström (SLU Umeå) for generously providing the images used in this paper.

1 Introduction

If one of your neighbors decides to purchase a good, does that affect your decision to buy the same good? In this paper, I study this question using household level data collected in a novel way. I examine a large set of aerial photos to infer household consumption of a good, specifically a trampoline. To collect a sufficiently large data set in a cost-effective manner, I automate the data collection process using a state-of-the-art neural network trained to identify trampolines. Neural networks are a class of non-linear classifiers capable of advanced pattern recognition. They have revolutionized (Krizhevsky et al., 2012) algorithmic image classification and are now rivaling human performance in some real-world classification tasks (Haenssle et al., 2018). This study provides an example of how neural networks can be applied to accurately collect economically relevant data from an image-based source on a scale which was infeasible just a few years ago.

A common motivation for the study of peer effects [1] is the discovery of social multipliers. For example, if low-performing pupils impact their classmates negatively, addressing the needs of these students may generate additional benefits by improving the outcomes of their peers (Lavy et al., 2012). Take-up of government programs can increase without the need for expensive information campaigns provided that information about the reform is transmitted between peers (Dahl et al., 2014).

In the consumption literature, the bulk of empirical research on peer effects concerns total household consumption expenditures and it's life-cycle pattern. While total consumption expenditures is arguably the most relevant outcome in examining the macro-economic implications of social multipliers, studying demand for specific goods is necessary to gain an understanding of the nature of peer effects. As Moretti (2011) notes, the limited empirical work on spillovers in the consumption of specific goods is likely caused by the lack of data. The ideal data set does not only need to record individual-level purchases of a specific good, it must also contain information on the relevant network and the purchases made by peers. Additionally, the data set needs to be sufficiently large to give the researcher the degrees necessary to overcome the econometric obstacles to identifying peer effects.

In this study, I use a novel technique to construct a unique large-scale data set on household-level trampoline ownership. By observing the exact location of each residence, I am able to form a complete picture of the neighborhood consumption profile. By using multiple consecutive observations of the same property, I am able to identify the time of purchase as well as focus on within-household change to implicitly control for many confounding factors. To this I add data

[1] Here, I use "peer effect" and "network effect" interchangeably

on the universe of real estate transactions, including the exact date when a specific property was bought/sold. My main identification approach relies on conditionally exogenous change in household exposure to neighboring trampolines that occur when an old neighbor moves out and a new one moves in. By comparing incumbent households living in the area, where one is "treated" with a new neighboring trampoline and one is not I am able to identify the effect of an additional neighboring trampoline on the purchase decision of the household. My results indicate that a neighborhood effect is present, as additional neighboring trampoline increase the probability of subsequent ownership by about 0.7 percentage points. The effect likely has a modest impact on the aggregate demand for this particular product. My estimates are similar in size to what has been found by Kuhn et al. (2011) and Grinblatt et al. (2008).

Conspicuous consumption and keeping-up-with-the-Joneses-effects are perhaps the most intuitively appealing explanations of peer effects. De Giorgi et al. (2016) find evidence of peer effects between co-workers which is broadly consistent with status-seeking explanations. However, when looking into expenditures for specific categories of goods, the authors find no evidence in favor of household demand shifting in favor of goods consumed by, or visible to, peers. For the latter exercise, the authors rely on a small sample consumer expenditure survey. They note that this particular exercise is likely underpowered.

There are complementary explanations of peer effects that directly relate to the characteristics of the individual good, such as social influence and information transmission (Dahl et al., 2014; Grinblatt et al., 2008; Moretti, 2011) as well as network externalities (Gilchrist and Sands, 2016). Grinblatt et al. (2008) studies neighborhood effects in the context of Finnish car purchases. They find that neighborhood effects are present, but that information transmission between neighbors is a more likely explanation than status-seeking. Kuhn et al. (2011) studies effects on neighbors of Dutch lottery winners. Consistent with Grinblatt et al. (2008), they find that spillovers are limited to car purchases.

This study complements previous work by applying a new data collection method in order to study a new type of product. Several studies have focused on cars purchases. From a macro perspective, cars represent an important consumer durable. As such, any social multiplier may have significant impact on the macro economy. However, if we want to deepen our general understanding of peer effects in consumption, we have to broaden our focus as a car is an outlier in terms of e.g. consumer expenditures. While a car purchase is likely to be preceded by months of deliberation and information gathering, Grinblatt et al. (2008) finds that neighborhood effects are limited to 10 days after a purchase. The authors speculate that a neighbor may "tip the scale" in favor of a purchase by providing word-of-mouth information about e.g. pricing at local dealerships. For a relatively inexpensive and homogeneous good like a trampoline, a purchase is likely more spontaneous and information transmission between neighbors may be less relevant. While the choice to study trampolines may seem peculiar, I argue that a trampoline has a set of unique characteristics that makes it suitable to study. First, it is a visible durable good designed for outdoor use.[2] For safety reasons, it is unlikely to be placed next to walls or trees, thus making it easy to observe in an aerial photograph. This characteristic is obviously essential to my data collection. Second, a trampoline is highly visible to neighbors, meaning that neighborhood effects can arise even under limited social interaction between neighbors. It is also durable, meaning that neighbors are exposed to it for a long time.[3] Third, it is a product used by and marketed to children. I conjecture that children are more likely to have local social network than adults, increasing the likelihood that the neighborhood is a relevant network to study.

As with any study concerning a specific good, there are idiosyncrasies that may limit the external validity of my findings. Regarding the underlying mechanisms, we have to consider the implications of a good that is marketed to children. To an adult, a trampoline is arguably not desirable from a status-seeking perspective. However, to the extent that the preferences of the child enter into household decision-making, through parental altruism or as outcomes of a bargaining process, status-seeking behaviour may manifest. My prior is that a neighborhood effect, if present, should be positive. However, one could consider a trampoline an (impure) public good, meaning that a neighboring trampoline may act as a substitute to a self-owned. Ultimately, my reduced form estimates are only informative about the net marginal effect.

Another contribution of this paper is to serve as an example of how neural networks can be used to tap into a previously inaccessible mountain of data. An increasing share of all data generated by humans comes in the form of images, be it aerial photos, physical documents or images posted on social media. Up until very recently, interpreting and compiling images was largely a labor-intensive manual task. This has prevented all but the most well-funded researchers to identify and transform images to a form that allows for conventional analysis. Neural networks are now easy to access and apply, expensive hardware and a degree in computer science are no longer prerequisites. So far, there are few examples of neural networks and computer vision applied to images within economics. Henderson et al. (2012) uses night-time illumination to improve measures of growth in developing countries. However, information is inferred directly from individual pixels and no interpretation is necessary. Gebru et al. (2017) uses Google Street View photos of parked cars and assign them a make and model using a neural

[2] The most popular trampolines listed on the leading Swedish price-comparison site have a diameter of 3-4 meters and a price of 2,000 to 10,000 SEK (240 to 1,200 USD).

[3] Kuhn et al. (2011) interpret their findings as visibility and durability being crucial for generating neighborhood spillovers.

network. They then use the composition of vehicles to predict the socio-economic composition of a neighborhood. They find that their derived measures of e.g. median income correlate well with survey data, but with a clear advantage in terms of cost and speed of measurement.

The paper is organized as follows. The next section presents the data along with the training and performance of the neural network (further information is presented in appendix A.1). Section 3 discusses my empirical framework and identification. Section 4 presents my results and section 5 concludes.

2 Data

My primary data source is a set of aerial photos of Swedish neighborhoods. The photos are produced by the Swedish mapping, cadastral and land registration authority (Lantmäteriet). The images cover the major urban centers and surrounding areas. My specific subset of images are selected by virtue of containing a large number of single-family homes, which is my main unit of observation (table 7 presents the distribution of property lot by year observations at the municipal level). The photos are taken during 2006 until 2017. The observation weighted mean time between consecutive photos of a property is 2.5 years and is on average photographed 4.2 times during my sample period. By slicing the raw images into smaller "chips", each containing a single property, I am able to apply the neural network to infer the presence of a trampoline on any given property (see section A.1 for details on image prepossessing). The training and performance of the neural network is described in the next section. Ultimately, the network produces a binary vector where 1 indicates the presence of a trampoline on a particular property at the time the photo was taken. To able to map trampoline ownership to a specific household, I exclude other forms of housing tenure from my sample. Around 60 percent of Swedish households with young children reside in a single-family home. While many reside in multi-dwelling buildings, these are typically located in urban centers, which effectively makes trampoline ownership infeasible.

To the trampoline ownership data I add data on real estate transactions, including the coordinates of the property, price and date of sale as well as the size of the home. The coordinates and date of sale allows me to match the transaction data to observations on trampoline ownership for a specific house. By comparing the transaction date to the photo date, I infer whether or not a new neighbor has moved in just prior to the photo date. If, between two consecutive photos, no transaction is recorded, I infer that the same household resides in that home in both periods. This allows me to include household fixed effects in my model.

In addition to property level data, I have data based on two larger spatial groupings, grid

Table 1: Summary statistics

	Obs.	Mean	SD	Min	Max
Houses per district	2,120	257	203	1	889
District population (2017)	2,120	1,771	417	720	3,297
Share ages 0-15 years	2,120	.206	.0566	0	0
Grid square population	45,906	117	138	0	1,768
Share ages 0-15 years	45,503	.218	.118	0.00	1
Grid square median income (2010 SEK)	45,882	275,070	93,552	0	506,652
Trampoline (T) present	2,148,452	.126	.331	0	1
Neighbors (N) with T	2,148,452	1.22	1.27	0.00	10
Nbrs moved in, past 60 days	2,148,452	.0677	.309	0	10
Nbrs moved in, past 100 days	2,148,452	.114	.4	0	10
Nbrs moved in, past 200 days	2,148,452	.231	.572	0	10
Nbrhd ΔT due to new N, 60 days*	1,537,223	-.00122	.0992	-3	2
Nbrhd ΔT due to new N, 100 days*	1,537,223	-.000846	.129	-3	2
Nbrhd ΔT due to new N, 200 days*	1,537,219	.00148	.184	-3	3
Transition from T to no T ($t-1$ to t)	1,537,095	.0535	.225	0	1
Transition from no T to T ($t-1$ to t)	1,537,095	.0653	.247	0	1
No. of properties	544,320	.	.	.	.
Property area (sqm.)	544,320	804	536	35	5,176
No. of households	611,357	.	.	.	.
Photos per household	.	3.51	1.23	1	6
House price (2010 SEK)	95,343	3,027,381	1,979,269	4,326	45,085,713
House area (sqm.)	89,986	129	38.7	10	290

*Conditional on same owner in $t-1$ and t

square and district. I add data data on demographics and median income at a spatial grid level. The grid is the highest spatial resolution for household data available from Statistics Sweden[4]. In urban areas, the grid has a resolution of 250 by 250 meters and in rural areas the resolution is 1 by 1 km. This data is available from 2014 until 2017, which is the later part of my main sample period (2006-2017). My main spatial grouping are districts. Districts[5] are constructed by Statistics Sweden and subdivides Sweden into about 6,000 areas, each containing 700 to 3,300 residents. Rather than aiming for a fixed population size, the district borders account for local geographic features such as streets, waterways, railroads and other infrastructure. This feature makes the districts ideal as a control for local economic conditions, advertising by local retailers and other supply side factors. My sample covers 2,120 districts and each district contains on average 257 single-family homes (see table 1 for summary statistics).

[4] Coordinate level data on e.g. income is typically not provided to researchers.
[5] Officially called Demographic Statistics Zone (DeSO)

2.1 Training the neural network

The neural network used in this study is called Inception-Resnet (Szegedy et al., 2016).[6] This publicly available neural network has demonstrated state-of-the-art performance on the benchmark ImageNet (Deng et al., 2009) dataset. In addition to excellent performance, it is less memory intensive than competing networks, making it feasible for consumer grade hardware.[7]

After preprocessing the aerial photos (see section A.1), I am left with a set of about 2,148,452 images, each containing a single property lot with a single-family home. To train the network to recognize a trampoline, I create a training data set by manually labeling a subsample of 22,435 images as being a member of either of the two classes, "Trampoline" (2,375 images) or "No trampoline" (20,060 images). Images of trampolines are underrepresented due to the simple fact that trampolines are fairly rare. The frequency imbalance between the two classes poses a problem for training because the iterative updating of the network weights can easily take the loss function to a local minimum where every image is predicted to be of the majority class without actually learning anything about the features of the images. To induce learning, I manually oversample the underrepresented class in the training and validation data to achieve a 50-50 balance of images containing trampolines and images without (Buda et al., 2018).

In the initial stages of training, I manually inspected all erroneous predictions made by the model in order to identify potential weaknesses. I initially found that for type 1 errors. i.e. finding a mat in an image labeled "No trampoline", there was often, upon review, actually a trampoline present in the image. I had simply missed it during labeling due to e.g. partial occlusion, low contrast, placement or, most importantly, human error. While this study is not a formal comparison between manual and automated image classification, this provides some anecdotal evidence that the accuracy of manual classification in this seemingly trivial task is being challenged by neural networks.

While it is reassuring that the model can identify even elusive trampolines, it is precarious to let the model influence the labeling of training data. The risk is that any flaws in the model is reinforced during training and that the "ground truth" that should be represented by the label is compromised. This makes it impossible to assess true model performance. Rather than simply trusting the model when identifying an incorrectly labeled image, I set up a script that fetches an image of the same property using the Google Maps Static API.[8] This allows me to

[6] As a side note, I initially tried to train an older network called VGG16, the 2014 winner of ILSVRC, with poor results. I was only able to achieve high accuracy by overfitting the training data, suggesting that the network was not "deep" enough to capture robust representations of the images.

[7] I perform all training and classification on a standard desktop computer equipped with an Nvidia Geforce GTX 1080ti GPU using the open source Keras Python library with a Tensorflow backend.

[8] Using only Google Maps data for this study was not possible at the time of writing for several reasons.

Figure 1: Examples of properties where the trampoline is difficult to make out in the original image (columns 1 and 3). In the high resolution Google-sourced image of the same property (to the right of the original), it is easy to verify the presence of a trampoline.

get two independently sourced images of the same property, which is helpful in borderline-cases. Additionally, the Google Maps image is often of higher resolution and taken from a different angel. In ambiguous cases where the model and I disagree on the correct label, I examine the corresponding Google Maps image and use it as a tie-breaker. Figure 1 present a few examples of the visual differences between the two image sources.

With images labeled as correctly as possible, I then randomize the data into an 80-10-10 training-validation-testing split. The validation set is used for model selection while the testing set is used to evaluate final model performance. During training, I use the Adam optimization algorithm and an initial learning rate of 0.001. When validation accuracy stops improving, the learning rate is reduced by half and so on until the accuracy converges. I set the maximum number of epochs to 30 with each epoch consisting of a complete pass of all training images, but the model typically converges faster. To get the most mileage out of my relatively scarce trampoline images, I augment the training and validation data by applying random shifts and flips around each axis and random 90-degree rotations. This has been shown to increase model

Table 2: Model performance, probability threshold = 0.5

		Label ($N = 1,819$) Trampoline	No trampoline	
Prediction	Trampoline	241	34	Precision = 87.6 %
	No trampoline	7	1,537	
		Recall = 97.2%	FPR = 2.8 %	F1-score = 92.2%

robustness (Vasconcelos and Vasconcelos, 2017).

Table 2 summarizes final model performance. While the model performs well in terms of recall (the share of labeled trampolines found), the precision (share of predicted trampolines that are correct) is somewhat lacking. The reason for this is that trampolines are rare, and there are many negatives that the model has to get right. Regarding false positives, I find that the network is sometimes fooled by similar structures, such as a round jacuzzi or a gazebo. Since I rely on within-household variation to identify effects, a prediction error that is due to some permanent feature of the property is likely to be repeated over time meaning that the error will not contribute to the identifying variation. The final output of the model is a vector of probabilities, where probabilities close to 1 represent a high degree of certainty that a trampoline is present. As figure 2 illustrates, there is a trade-off between precision (type 1-errors) and recall (type 2-errors). Put differently, I can make the model find fewer false trampolines by increasing the probability threshold, but in doing so I will also increase the likelihood that the model misses a few actual trampolines. As the figure shows, increasing precision entails a substantial drop in recall. In my predictions, I use the default 0.5 probability threshold to infer trampoline ownership.

It is important to relate performance to alternative approaches. Typically, consumption data is obtained through surveys which are known to be less than perfectly accurate. One Swedish study found underreporting of car purchases by up to 30 percent (Carroll et al., 2015). Given that car expenditures make up a large fraction of household consumption, a survey of trampoline expenditures is likely to be burdened by even higher error rates. Using credit card transaction data to measure consumption has grown increasingly popular (e.g. Agarwal et al., 2017a). While this data addresses reporting error, it cannot cover the full universe of transactions and researchers are often limited to customers of a specific financial institution. This is particularly problematic for the study of network effects, since the researcher is not able to get a complete picture of the relevant network.

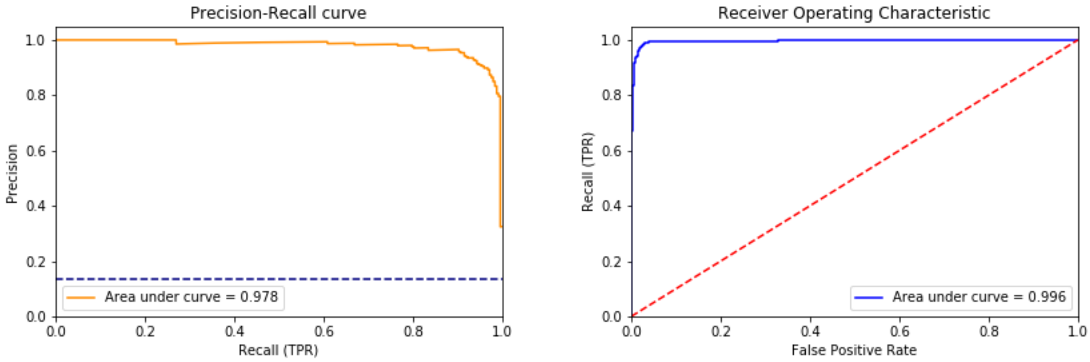

Figure 2: The two panels illustrate the trade offs between precision and recall (left) and between precision and false positives (right) by plotting the respective statistics for a number of different probability thresholds. The dashed line on the left is the maximum precision possible a model with complete overlap between the two classes (a coin-toss prediction) and is given by the ratio of *Trampoline* to *No trampoline* images in the test set (13 %). To the right, a coin-toss prediction will be located somewhere along the dashed line.

Figure 3 shows mean trampoline ownership rates as predicted by the neural network. To avoid confounding a time trend with sample selection due to the rotation of houses being photographed each year, I plot the trend for the six most common photo-year sequences. All of them depict an upward trend in trampoline ownership in the sample, going from under 10 to around 15 percent. The increase in popularity is corroborated by data from Google Trends. As shown by Figure 4, the number of searches for "trampoline" has increased steadily since 2012. There is a strong seasonal component, with Sweden's cold winters freezing the public's interest between October and March. Mentions of trampolines in news media has declined since the peak in 2010.[9]

3 Empirical framework

My goal is to estimate the effect of neighbors' decision to acquire a trampoline on the ownership status of the household. This outcome-on-outcome effect belongs to what Manski (1993) refers to as an endogenous social effect or what is now commonly known as a peer effect. In the peer effects literature, a large body of work concerns identification. The main identification issue stem from the fact that peer groups are rarely formed by chance alone. Formation is often heavily influenced by sorting and self-selection. The formation process therefore tends to produce groups whose members have similar characteristics and preferences and are subject to similar shocks. Further complicating the interpretation of estimates is what Manski (1993) called "the reflection problem". In a simple contemporaneous outcome-on-outcome regression, the direction

[9] Data from Retriever Media Archive. Between 2000 to 2004, mentions of trampoline are much less frequent. The recent decline is likely due to the fading novelty of trampolines as a common household product. In addition, EU regulation introduced in 2014 made safety enclosures a mandatory feature of large trampolines. This has likely decreased injuries, which used to be a common theme in media coverage.

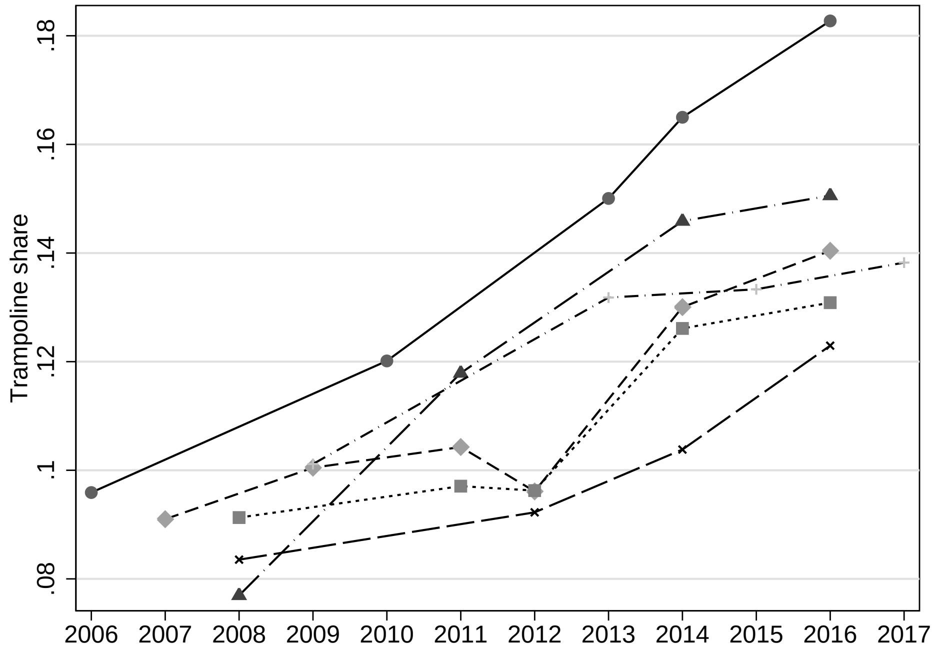

Figure 3: Mean share of homes in the sample with a trampoline, as predicted by the neural network. Lines represent the most common photo-year sequences.

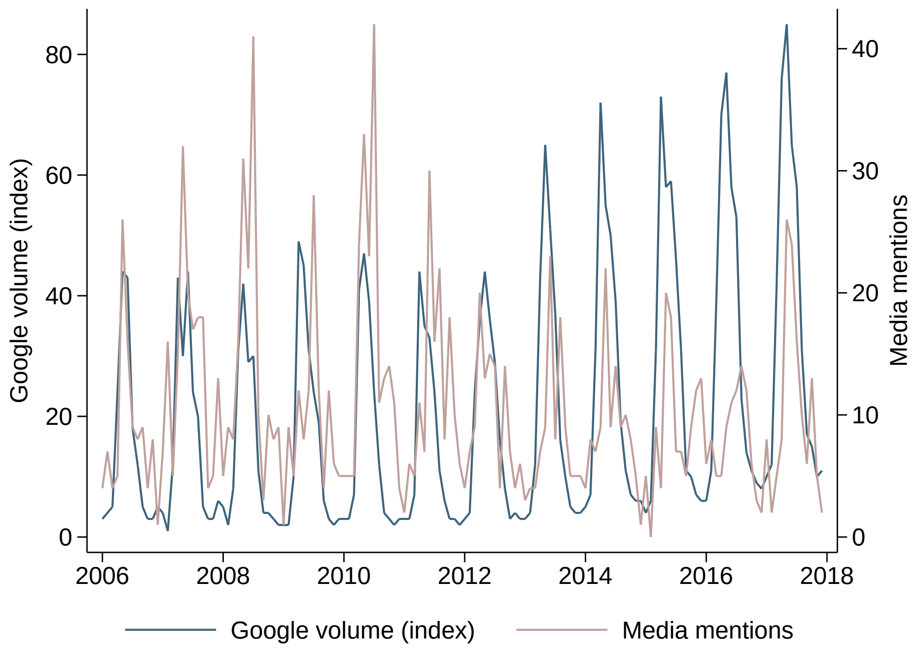

Figure 4: Trends in Google search volumes and media mentions for the keyword *studsmatta* (*trampoline*). Monthly observations.

of causality cannot be inferred unless the researcher can distinguish action from reaction. For example, is a student performing better in school because their peers have improved, or have the peers improved because of the subject? Due to this simultaneity of outcomes, we cannot interpret the estimated coefficient as a social interaction effect but have to rely on e.g. instruments for the outcome of peers.

My point of departure is the model explored by Bramoullé et al. (2009) in their paper on identification of peer effects in social networks:

$$y_{it} = \alpha_i + \beta W_i \mathbf{y}_t + \gamma \mathbf{X}_{it} + \delta W_i \mathbf{X}_t + \varepsilon_{it} \tag{1}$$

Here, the outcome observed for household i in period t is a function of $\mathbf{y}_t$, the $n \times 1$ vector of contemporary outcomes for all other households. Equation (1) is inspired by the spatial econometrics literature which commonly makes use of a spatial weight matrix $\mathbf{W}$ to capture the spatial arrangement of units. $\mathbf{y}_t$ is multiplied by W_i, the i^{th} row in the $n \times n$ spatial weight matrix $\mathbf{W}$. The spatial weight matrix is partially given by the local geographic arrangement of homes, but some further modelling is required. In the next section, I explain how I construct $\mathbf{W}$. β represents the *endogenous* social interaction effect that is the focus of this paper, i.e. the effect of my neighbor getting a trampoline on my own purchase decision. The specification in Bramoullé et al. (2009) also includes the *exogenous* social effect δ, i.e. the effect of my neighbors exogenous characteristics on my own purchase decision. While there is not much intuition for this type of effect for this particular good, one could imagine that my purchase decision is affected by the ages and number of children that live next door. In my regressions, I control for demographic factors.

In the absence of correlated effects, i.e. innovations to ε or $\mathbf{X}$ that are common to the neighborhood, Bramoullé et al. (2009) shows that non-overlapping peer groups, formalized as the linear independence of $\mathbf{I}$, $\mathbf{W}$ and $\mathbf{W}^2$ is sufficient to be able to identify $(\alpha, \beta, \gamma, \delta)$. Another way of stating this condition is that my neighbors have to have neighbors who are not my neighbors, meaning that we can form an *intransitive triad* of households. This provides a set of exclusion restrictions, as any shock to the neighbors of my neighbors may only affect me through my neighbors. Intransitive triads have been used to identify network effects in empirical research (De Giorgi et al., 2016). While the presence of intransitive triads is ultimately an empirical question, it is hard to infer their existence from from geographic arrangement alone. Therefore, I opt to use a source of conditionally exogenous variation in $W_i \mathbf{y}_t$ that directly affects my outcome. Specifically, I use changes in neighboring trampoline composition caused by new neighbors moving in.

As pointed out by Angrist (2014), even if β is identified in an econometric sense, it can only be interpreted as a social interaction effect if the error term is not correlated within neighborhoods. To address this, Bramoullé et al. (2009) proposes the inclusion of a neighborhood fixed effect, α_i, which absorbs the effect of time-invariant unobservables. As my data includes the exact location of each household, every household forms the basis for a neighborhood. The household fixed effect α_i therefore subsumes a neighborhood fixed effect.[10] To further control for area-specific trends, I include a full set of district by year fixed effects.

According to Angrist (2014), the most compelling studies of peer effects uses variation in peer group composition that is orthogonal to baseline characteristics. In the spirit of this ideal, I exploit variation stemming from new neighbors moving in as a source of plausibly exogenous change in household trampoline exposure (from the perspective of incumbent neighbors). For a discussion of the identifying assumptions behind this approach, see section 4.2. Under the assumption that a new neighbor moving in constitutes a (conditionally) exogenous change in the trampoline profile of a neighborhood (see section 4.2), I use the following specification to estimate the neighborhood trampoline effect:

$$y_{it} = \alpha_i + \beta W_i D_{it-1} \mathbf{\Delta y}_{t-1} + \gamma \mathbf{X}_{it} + \varepsilon_{it} \tag{2}$$

The key difference between equation (2) and (1) lies in how I specify treatment, $D_{it-1}\mathbf{\Delta y}_{t-1}$. The first term, D_{it-1}, is the i^{th} row in the move matrix $\mathbf{D}_{t-1}$. $\mathbf{D}$ is a binary counterpart of $\mathbf{W}$ where an element d_{ij} takes on a value of 1 if the j^{th} neighbor of i moved-in just before the period $t-1$ photo was taken and 0 otherwise. The second term, $\mathbf{\Delta y}_{t-1}$, is motivated by my assumption that for a social interaction effect to arise, we need incumbent neighbors' exposure to trampolines to change. I argue that this is true for both a learning and status-seeking channel. Consequently, if a new neighbor does not contribute to a change in the neighborhood trampoline profile, there is no treatment. This occurs if a trampoline (no-trampoline) household moves into a home where a trampoline (no-trampoline) used to live. Therefore, I specify treatment in first differences.

The third important aspect of (2) concerns the issue of simultaneity and timing of the effect. As I conjecture that a new neighbor constitutes a conditionally exogenous change in the neighborhood trampoline profile (I discuss this further in section 4.2), specifying a contemporaneous effect of $D_{it}\mathbf{\Delta y}_t$ would implicitly assume a unidirectional effect in the sense that incumbent neighbors can be affected by new neighbors in the same period, but new neighbors are unaffected by the incumbent. By lagging treatment, I give incumbent neighbors time to react to the

[10] The two are not equivalent since the neighborhood persists through households moving in and out.

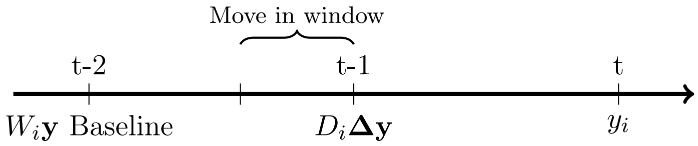

Figure 5: Timeline illustrating my empirical approach of using new neighbors to identify the neighborhood effect.

change in neighborhood trampoline profile. Hence, we measure the reaction of household i in period t to a change in the trampoline profile that occurs between $t-2$ and $t-1$. Grinblatt et al. (2008) presents a similar argument for dealing with the reflection problem. If treatment precedes the outcome, there is a limited scope for the effect to run in the "wrong" direction.

As a secondary approach, I estimate a specification similar to Grinblatt et al. (2008). By including outcomes for near and far neighbors simultaneously, the authors effectively difference-out the neighborhood fixed effect. While highly localized *unobservable* differences may still bias the estimates, the authors argue that since *observable* differences between near and far neighbors are small, it is unlikely that any unobservable differences can explain the result (in absence of a true neighborhood effect). In my framework, I implement this approach using the following specification:

$$y_{it} = \alpha_i + \sum_{k=1}^{K} \beta_k W_{ki} \mathbf{y}_{t-1} + \gamma \mathbf{X}_{it} + \varepsilon_{it} \tag{3}$$

By including K sets of neighbors separately, I am able to compare the effect of close (e.g. β_1) to more distant neighbors (e.g. β_2). Under the assumption that there are no unobservable differences between near and far neighbors related to trampoline demand, (3) identifies the neighborhood effects.

3.1 Modelling neighborhoods

The role of the spatial weight matrix $\mathbf{W}$ is to describe which households are to be considered neighbors in the sense that their respective trampoline ownership status is communicated (directly or indirectly). As noted by Anselin (2009), there is little in the way of general guidance in how $\mathbf{W}$ should be specified. Researchers commonly use geographic distance and contingency. Rather than starting from an arbitrary choice of $\mathbf{W}$, I opt for a more data-driven approach. Suppose that the relationship between the trampoline status of household i and its M closest neighbors can be described as:

$$y_{it} = \alpha_i + \sum_{j=1}^{M} \beta_j y_{jt} + \varepsilon_{it} \tag{4}$$

Where β_j represents the effect (in a purely correlational sense) that household j's status has on household i. To test the role of distance between neighbors in determining between household correlation, I apply the Least Absolute Shrikage and Selection Operator (LASSO) to (4). The LASSO is a commonly used model selection tool developed by Tibshirani (1996) that imposes a measure of regularization to standard OLS in order to prevent overfitting caused by superfluous variables. Formally, LASSO is defined as the following Langrangian optimization problem:

$$\beta_L = \underset{b_L}{\arg\min} \{ \frac{1}{N} ||\mathbf{y} - \alpha - \mathbf{Y} b_L||_2^2 + \lambda ||b_L||_1 \} \tag{5}$$

where α is an N-vector of household fixed effects, $\mathbf{Y}$ is the $N \times M$ matrix describing the trampoline status for the M closest neighbors to each household and b_L is an M vector of parameters to be estimated by the LASSO. The second term represents the regularization term, operating on the L1-norm of b_L (the sum of absolute values). The Lagrange multiplier λ is a penalization coefficient, determining the relative importance of minimizing $|b_L|$. Because the LASSO uses L1-regularization, the solution will force a subset of parameters in β_L to be exactly equal to zero. This is the model selection property of the LASSO that allows me to determine which neighbors are the most important predictors of a within-household change in trampoline status. In the estimation procedure, λ is a free parameter to be determined by the researcher. To explore the role of geographic distance in determining the size of a neighborhood, I use a range of values for lambda. While there are several methods for finding an optimal value[11] for λ, for this exercise I believe it is more informative to determine how the β_L evolves under different levels of regularization.

I consider the 50 closest neighbors to each household in my sample ($M = 50$).[12] Figure 6 presents the LASSO results. Two noteworthy patterns emerge. First, even within this highly local scope, distance seems to matter. For the largest penalty, only the closest two neighbors are selected. As λ decreases, neighbors 3 trough 9 are also selected. Second, as we lower the penalty further, distance seem to matter less. We notice that as λ drops below 1000, the selection of new neighbors seems more or less unrelated to distance. This could indicate that neighborhood effects become less important than e.g. effects common to a greater area at greater distances.

[11] This is typically done by minimizing prediction error or an information criteria.

[12] While the number 50 is arbitrary, what is critical is picking a number larger than the number of neighbors one could reasonably expect each household to interact with.

Given my interpretation of figure 6, I consider $M = 10$ as my baseline neighborhood definition as the 10 closest neighbors seem to be the most important in terms of model selection as well as their impact in the post-estimation OLS.

By setting $M = 10$, I standardize the number of non-zero elements w_{ij} in each row i of $\mathbf{W}$ to 10. However, as some houses straddle the border of the area covered by the aerial photography, one or more of their neighbors may fall outside the area of coverage. As a result, some homes will have less than 10 effective neighbors.[13] As for the value of each element w_{ij}, I consider both the the inverse distance weighted average ($w_{ij} = \frac{\omega_{ij}}{\sum^{M} \omega_i}$) as well as simply counting the number of neighboring trampolines ($w_{ij} = 1$). On average, the centroid-to-centroid distance to the 10^{th} neighbor is 75 meters and 195 meters to the 50^{th} (Figure 12 describes the distance to the 15 closest neighbors). Interestingly, Grinblatt et al. (2008) also uses the 10 nearest neighbors as their main definition of a neighborhood.[14] Other studies have also found that neighborhood effects seem to transmit only over short distances. Agarwal et al. (2017b) only finds effect among residents in the same residential building and Kuhn et al. (2011) find that effects are limited to neighbors "within two doors".

4 Results

I begin this section by presenting results of my estimation of the Grinblatt et al. (2008) approach using equation (3). I then discuss and test the validity of using new neighbors as a source of exogenous variation in the neighborhood trampoline profile. Next, I present results using the variation stemming from new neighbors moving in. Finally, I put the effect size into context by performing some simple comparisons to aggregate demand.

4.1 Neighborhood fixed effects

My first set of results presents my version of the approach by Grinblatt et al. (2008), see equation 3. To exploit the full variation in the data, I do not limit the explanatory variable to new neighbors. To account for simultaneity, I estimate (3) with both a contemporaneous and lagged treatment. Table 3 presents the results. In columns 1 and 2, I specify treatment as the share of the 10 nearest neighbors with a trampoline, weighted by inverse distance. In addition to a household fixed effect, I control for correlated effects and common shocks by adding a full set of district by year fixed effects well. With weights all normalized to sum to one, the estimates in columns 1 and 2 represent the effect of all 10 neighbors acquiring a trampoline. While

[13] 93 percent of all home have all 10 neighbors covered. 99 percent have 7 or more.
[14] The paper does not provide a clear motivation for the choice of 10 specifically.

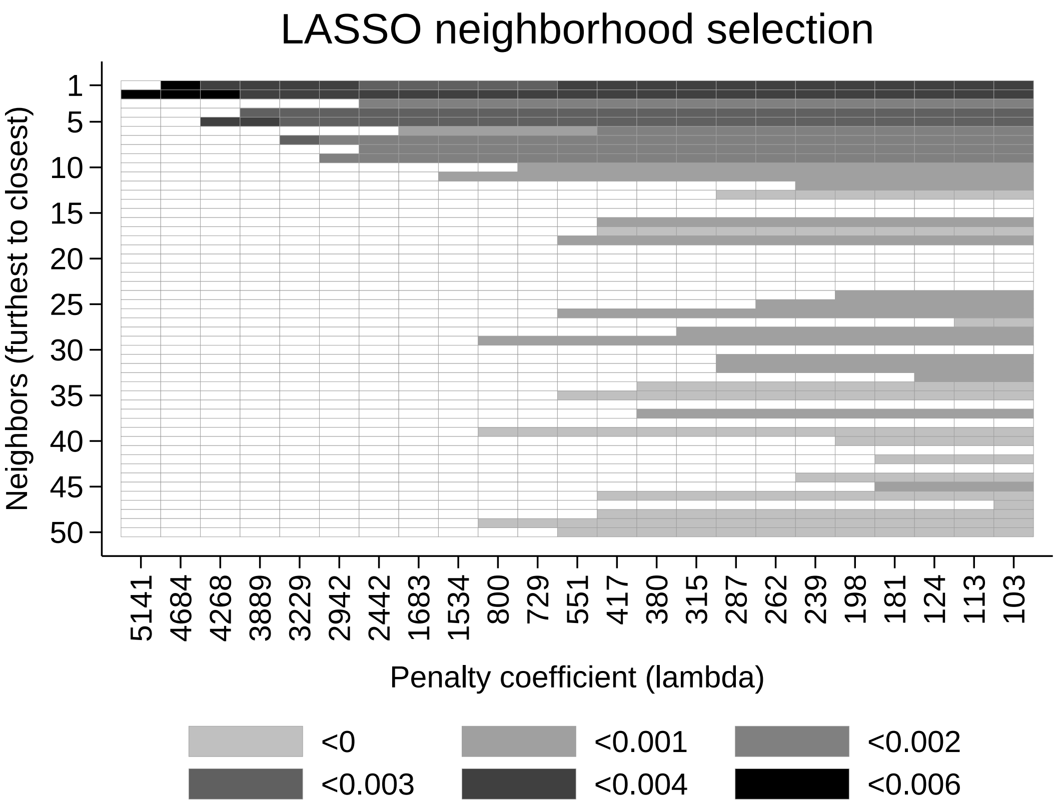

Figure 6: The figure depicts the subset of neighbors selected by estimating equation (5) using a range of different values for the penalization coefficient λ. The values used is the default sequence in the Stata package LASSOPACK. White areas indicates that the corresponding parameter is set to zero. Shaded areas indicate the value of the corresponding parameter estimate in an OLS regression on the LASSO-selected subset of neighbors.

contemporary effect is quite large (column 1), when lagging treatment the effect is reduced to about 1 percentage point off a sample average of about 13 percent (column 2). The reduced effect size is consistent with simultaneity bias inflating the contemporary estimates.

I also include neighbors 1-5, 6-10 and 11-15 separately (columns 3 and 4). For marginal effects to be easily comparable across groups, I define treatment simply as the number of trampolines observed among each of the three groups. We note that even at this very local scale, the marginal effect of an additional trampoline rapidly diminishes.[15] While the fact that the effect diminishes even across very short distances could be attributed to a true neighborhood effect, we cannot rule out unobserved heterogeneity due to e.g. highly localized sorting within districts as a source of bias. However, I believe that this is an unlikely explanation. First, the mean distance between the 5^{th} and the 15^{th} neighbor is only 40 meters. To get a better idea of how similarity between households change over this short distance, I calculate the mean absolute deviation between each household and a group of neighbors J according to

$$M(X)_J = \frac{1}{N_J} \sum_{i}^{N} \sum_{j=1}^{N_J} |x_j - x_i| \tag{6}$$

By comparing $M(X)$ and raw means between groups of neighbors, we can quantify the differences. The major caveat is that my data at the household level is limited to just a few variables. The comparisons are presented in table 4. For house price and house area, I can only observe houses who are sold at some point during 2005 to 2017. We note that while all three variables have a higher MAD when comparing households to neighbors 11-15 as opposed to neighbors 1-5, the bulk of deviations to house i are present already among neighbors 1-5 (column 1 of table 4). Relative to the initial deviation, the increase is negligible. In the bottom panel, we also note that the raw means are almost equal between the two sets of neighbors. Finally, I would like to direct the reader to Table 3 in Grinblatt et al. (2008). They find that even over greater distances, neighbors remain broadly similar in terms of household-level demographics and income. Given the similarities between the Finnish and Swedish context, I would likely find similar patterns. To conclude, the data is not consistent with the idea of neighborhoods as consisting of small enclaves that are internally homogeneous and suggests that the results in table 3 cannot be explained by unobserved local heterogeneity.

[15] The difference between 1-5 and 11-15 is statistically significant in column 3 but not in column 4 (p-value = 0.12)

Table 3: Baseline regressions

Weighted share in t	3.857*** (0.269)			
Weighted share in $t-1$		1.524*** (0.317)		
Number of trampolines in t				
Neighbors 1-5			0.491*** (0.0425)	
Neighbors 6-10			0.169*** (0.0416)	
Neighbors 11-15			-0.00954 (0.0413)	
Number of trampolines in $t-1$				
Neighbors 1-5				0.131** (0.0504)
Neighbors 6-10				0.0847* (0.0494)
Neighbors 11-15				0.0204 (0.0495)
Observations	2,097,438	1,445,185	2,001,648	1,301,920
Households	560,522	468,796	570,393	441,002
R-squared	0.542	0.630	0.579	0.665
Outcome mean	12.58%	12.91%	12.65%	12.65%
Fixed effect	District × Year	District × Year	District × Year	District × Year

The table presents estimates of equation (5). All coefficients are expressed as percentage points. Standard errors clustered on households. *** p<0.01, ** p<0.05, * p<0.1

Table 4: Differences in MADs and means

	$M(X)_{1-5}$ (SE)	$M(X)_{11-15}$ (SE)	|Diff|
House area	24.16 (0.06)	27.86 (0.06)	3.70
House price	1,111,911 (2,620)	1,187,357 (2,806)	75,446
Property area	245.42 (0.63)	322.77 (0.70)	77.35
	Mean_{1-5}	Mean_{11-15}	|Diff|
House area	127.62	127.47	0.14
House price	3,170,214	3,153,883	16,331
Property area	852.14	847.55	4.59

4.2 New neighbors

Do new neighbors constitute a source of exogenous variation of the neighborhood[16] trampoline composition? This question merits some initial discussion and definitions. Because new neighbors self-select into homes and/or areas of residence, it would be naive to expect that the trampoline status of a new neighbor is orthogonal to the baseline trampoline profile and preferences of the neighborhood. For example, parents with young children likely sort into neighborhoods with a similar demographic profile. However, the notion that a new neighbor can provide a *conditionally* exogenous source of variation in the neighborhood trampoline profile is plausible. To clarify this idea, consider the following example: The Smith's are looking to buy a new home, to which they will be bringing their trampoline. While the Smith's preferences over the type of neighborhood and house they would prefer to live in, they cannot pick any home. They are limited to the homes currently available on the market. The week after they purchase their new home, a similar home in the same area is put on the market. The Smith's would have been just as happy in this new home, but due to the quasi-random timing of the supply of homes in the market, the trampoline-less Jones family moves in instead. From the perspective of incumbent neighbors, whether the Smith's or the Jones' become their new neighbors is as good as random.

Using a classic diff-in-diff framework, the identifying assumption is that the neighbors of the Jones' represent the counter-factual within household change in trampoline ownership that the Smith's neighbors would have experienced. We could imagine a world where people engage in highly localized sorting by some observable set of characteristics related to trampoline preferences that is not captured by a household fixed effect. This would violate the parallel trends

[16] In the following, neighborhood is defined as the 10 closest houses to each individual household unless otherwise specified.

assumption and bias my estimates. I present a few arguments against this hypothesis. First, I test if the prior trampoline profile can predict the trampoline status of new neighbors conditional on pre-determined area and property characteristics. Second, I visually inspect the pre-trends for household that I can observe for several periods prior to assignment into treatment assignment. Finally, in the next section I present some suggestive evidence that my results cannot be explained by highly localized sorting.

Another important qualification concerns the definition of a new neighbor. If neighborhood effects exist, it would be a mistake to assume that the newly moved-in Smiths and Jones are immune to these effects. As time goes by, their expected trampoline status is likely to converge to that of incumbent neighbors, be it because of shared characteristics, common shocks or an endogenous neighborhood effect. Therefore, establishing a post-move-in window during which new neighbors plausibly constitute an exogenous source of variation is crucial.

If we assume that trampoline ownership is sufficiently informative of the unobservable determinants of trampoline demand, we can perform a tentative test for the sorting of households on trampoline ownership using the following specification:

$$y_{itd} = \alpha + \beta W_i \mathbf{y}_{t-1} + \gamma' \mathbf{X}_{it-1} + \epsilon_{itd} \tag{7}$$

where y_{itd} is the trampoline status of household i who moved in between d and 7 days before the period t photo was taken. Here, I define $W_i \mathbf{y}_{t-1}$ as the weighted share of trampolines among i's neighbors in period $t-1$, i.e. before i moved in, and $\mathbf{X}_{it-1}$ is a vector of covariates. A statistically insignificant estimate of β suggests that $E[y_{itd}|W_i\mathbf{y}_{t-1}, \mathbf{X}_{it-1}] = E[y_{itd}|\mathbf{X}_{it-1}]$, i.e. that the trampoline status of the new neighbor is conditionally exogenous with respect to the neighborhood trampoline share in the previous period. The conditioning variables included in $\mathbf{X}_{it-1}$ are a full set of district by year fixed effects as well as decile indicators grid square median income, the share of the grid square population below 16 years of age, house and property area, mean area of neighboring properties and the sale price. A caveat is that my grid square level data is only available starting in 2013, so I am only able to include the later part of my sample.

The results from estimating equation (7) is presented by Figure 7. I run the regression for newly moved-in households, split by time living in the neighborhood before the photo was taken as well as for households who moved in more than two years ago (the red line). We note that the estimate of β for the newly moved-in sample approaches that for the incumbent sample, meaning that the prior trampoline profile becomes an increasingly strong predictor of trampoline status as time spent in the neighborhood increases. This is consistent with the idea that new neighbors adapt to their new neighborhood. Prior to 100 days spent in the neighborhood, we cannot reject

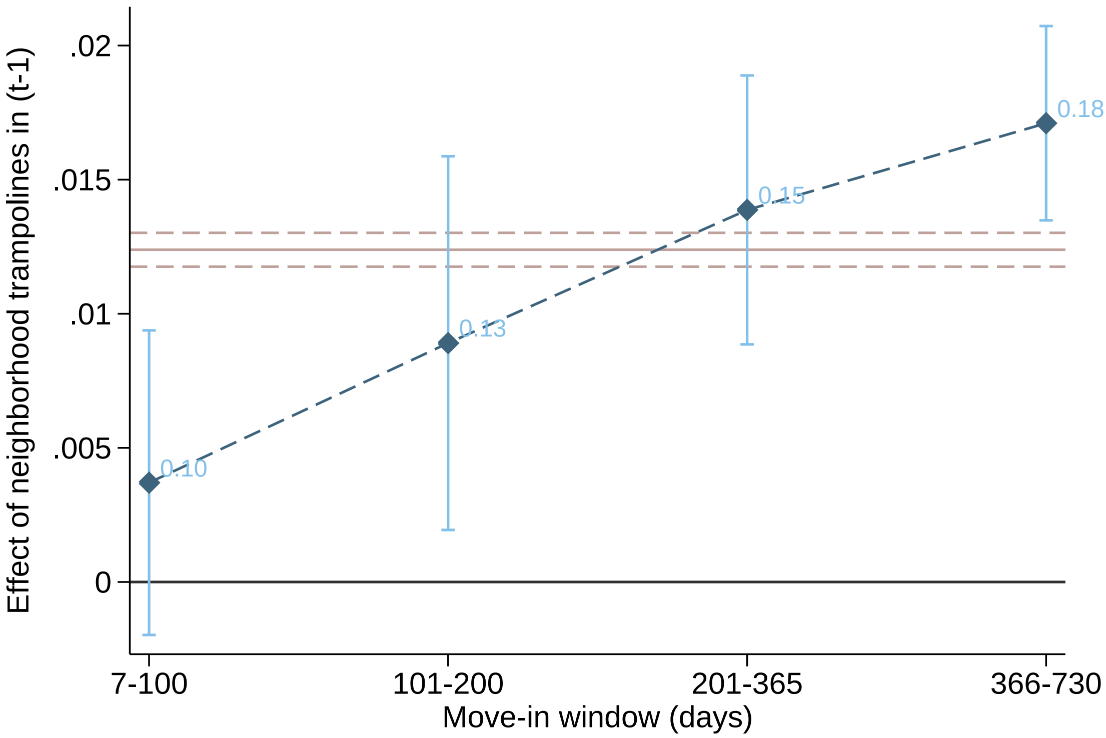

Figure 7: The figure presents estimates of β in equation (7) for a range of move-in windows prior to the period t photo was taken (in blue). The red is the estimate for households who moved in more than two years ago. Controls include a full set of district by year fixed effects as well as decile indicators for grid square median income, the share of the grid square population population below 16 years of age, house and property area, mean area of neighboring properties and the sale price. Labels indicate the mean of y_{itd}. Bars indicate 95-percent CI's.

the null of no effect at any conventional level. While the point estimate is still positive, it is important to keep in mind that in my main regressions, identification relies on within household changes to trampoline exposure. This implicitly controls for a broader set of confounding factors than I am able to using my limited set of pre-determined observables. The fact the the effect is eventually even greater than that for incumbent households is likely explained by the fact that household who move are more likely to have children and therefore have higher baseline probability of purchasing a trampoline. The marker labels in Figure 7 denote the mean of the outcome variable for the corresponding sample. As expected, the initial ownership rate among new neighbors is lower than the sample mean of about 13 percent but increases with time spent in the neighborhood.

To provide additional evidence that households who do not experience an increase the number of neighboring trampolines due to a new neighbor represent a plausible counter-factual to households that are treated with an increase in neighboring trampolines, I visually examine

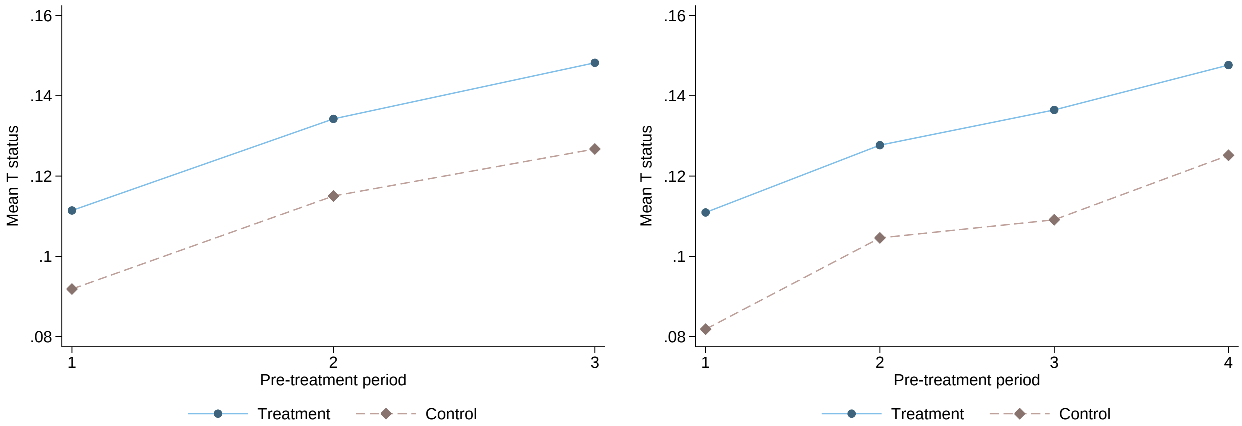

Figure 8: Pre-treatment trends for households I observe for 4 and 5 periods, respectively. All households receive a new neighbor within 200 days of the last photo date. The new neighbor(s) either contribute to an increase in neighboring trampolines (Treatment) or not (Control).

pre-trends for treated and non-treated households. The caveat here is that my panel is short and I am only able to observe pre-trends of any meaningful length for some households. I select the sample of households which I observe for 4 and 5 periods[17]. Within this sample, I focus on households who receive a new neighbor[18] just prior to the last time we observe them. I then split the sample into treatment and control groups, where treatment is defined as the new neighbor contributing to an increase in the number of neighboring trampolines and the control group experience either no change or a decrease due to the new neighbor. Figure 8 presents the pre-trends in terms of mean trampoline status. While short, the trends appear roughly parallel for the two groups.

Table 5 presents the results from my estimation of equation (2) for different move-in windows. As previously, the neighborhood includes the 10 closest neighbors and I define the variable of interest, $D_{it-1}\mathbf{\Delta y}_{t-1}$, as the change in the number of trampolines. Households who do not get a new neighbor in a given year have $D_{it-1}\mathbf{\Delta y}_{t-1}$, but are included in the regression as they help to identify district trends. While I have a lot of observations, the identifying variation stems from the subset of households who experience a change in trampoline status and/or a change in the neighborhood trampoline composition due to new neighbors moving in. I find a statistically significant effect when I limit the window to 60 days (column 1). While the standard error decreases for longer windows, so does the point estimate. This is consistent with longer windows providing less of an exogenous shock to incumbent neighbors, although the general lack of precision makes interpretation difficult. If the effects presented in table 3 represent the true magnitude of the effect, then the regressions in table 5 may be underpowered. However,

[17] The number of households observed for 6 periods is too small to be useful
[18] Moved in within 200 days

estimates are of the same magnitude as in table 3. The estimate in column 1 says that an additional neighboring trampoline increases the likelihood of a subsequent purchase by about 0.7 percentage points.

My prior expectation is that the neighborhood effect should be stronger for household who experience an increase in neighboring trampolines. While a decrease in neighboring trampolines could reduce the probability of acquiring one, the baseline probability of ownership is low, decreasing the scope for such an effect. It seems unlikely that a negative treatment would induce large number of household to sell their existing trampolines. This is partially confirmed by splitting the sample into positive and negative treatment (see the first two columns of table 6), as the point estimate is twice as large for the positive treatment group. In the rightmost two columns, I test how the size of the neighborhood affects my results. Focusing on the closest 5 or 15 neighbors produces estimates of similar size, but they are not statistically significant.

4.3 Is the effect economically significant?

To put the size of the estimates into context, it is important to recognize that households may choose to acquire or part with a trampoline for a myriad of reasons. To determine the relative effect size, I ask how large the neighborhood effect is in relation to total household demand. Consider a matrix form of equation (3):

$$\mathbf{y}_t = \alpha + \beta \mathbf{W} \mathbf{y}_{t-1} + \gamma \mathbf{X}_t + \varepsilon_t \tag{8}$$

To focus on the neighborhood effect, let us abstract from other covariates and assume that the only determinant of $\mathbf{y}_t$ is the interaction effect and the prior trampoline configuration $\mathbf{y}_{t-1}$ according to:

$$\mathbf{y}_t = \mathbf{y}_{t-1} + \beta \mathbf{W} \mathbf{y}_{t-1} + \varepsilon_t \tag{9}$$

$$\iff \Delta \mathbf{y}_t = \beta \mathbf{W} \mathbf{y}_{t-1} + \varepsilon_t \tag{10}$$

Let us consider households who do not own a trampoline in period 1. Of these, about 9 percent will purchase a trampoline in period 2. To what extent are these transactions explained by interaction with their trampoline-owning neighbors in period 1 as described by $\beta \mathbf{W} \mathbf{y}_{t-1}$? I take the trampoline configuration given by the first period we observe each house and allow each trampoline-owner to linearly increase the probability each of their 10 neighbors by $\beta = 0.01084$

Table 5: Results, new neighbors

Effect of $D_{it-1}\mathbf{\Delta y}_{t-1}$ using a move-in window of			
60 days	0.708* (0.398)		
100 days		0.311 (0.303)	
200 days			-0.0216 (0.206)
Observations	893,263	893,263	893,261
Households	385,627	385,627	385,626
R-squared	0.717	0.717	0.717
Outcome mean	12.56%	12.56%	12.56%
Fixed effect	District × Year	District × Year	District × Year

The table presents estimates of equation (2). The variable of interest, $\mathbf{\Delta y}_{t-1}$ is expressed as the number of trampolines. Standard errors clustered on households. *** p<0.01, ** p<0.05, * p<0.1

percentage points[19]. I find that exposure to neighboring trampolines explains at most 12 percent of all new trampolines in period 2. This represents an upper bound since my results suggest that the effect diminishes over time, meaning that only "new" owners actually give rise to the spillover. Additionally, 0.01084 is at the top of the estimate range. While these calculations are very much back-of-the-envelope, they suggest that the neighborhood effect has a modest impact on aggregate demand for this particular product.

5 Conclusion

TBD

[19] From a baseline of zero. 0.01084 corresponds to the top row of column 1 in table 6. I pick the largest estimate to get an upper bound of the aggregate effect. The baseline ownership rate in period 1 is 11.5 percent.

Table 6: Results, sensitivity checks

Effect of $D_{it-1}\mathbf{\Delta y}_{t-1}$ using a move-in window of				
60 days	1.084* (0.577)	0.494 (0.548)	0.478 (0.567)	0.485 (0.329)
100 days	0.742* (0.432)	0.0509 (0.435)	0.0874 (0.429)	0.384 (0.248)
200 days	0.264 (0.287)	-0.331 (0.310)	-0.253 (0.291)	0.113 (0.169)
Neighborhood size	10	10	5	15
Sample	$D_{it-1}\mathbf{\Delta y}_{t-1}$ $>= 0$	$D_{it-1}\mathbf{\Delta y}_{t-1}$ $<= 0$	All	All

The table presents estimates of equation (2), each cell is a separate regression. The variable of interest, $\mathbf{\Delta y}_{t-1}$ is expressed as the number of trampolines. All regressions include a full set of district $\times$ year fixed effects. Standard errors clustered on households. *** p<0.01, ** p<0.05, * p<0.1

References

Agarwal, S., J. B. Jensen, and F. Monte (2017a): "The Geography of Consumption," SSRN Scholarly Paper ID 3007488, Social Science Research Network, Rochester, NY.

Agarwal, S., W. Qian, and X. Zou (2017b): "Thy Neighbor's Misfortune: Peer Effect on Consumption," SSRN Scholarly Paper ID 2780764, Social Science Research Network, Rochester, NY.

Angrist, J. D. (2014): "The perils of peer effects," *Labour Economics*, 30, 98–108.

Anselin, L. (2009): "Spatial Hedonic Models," in *Palgrave Handbook of Econometrics*, ed. by T. C. Mills and K. Patterson, London: Palgrave Macmillan UK, 1243–1250.

Bramoullé, Y., H. Djebbari, and B. Fortin (2009): "Identification of peer effects through social networks," *Journal of Econometrics*, 150, 41–55.

Buda, M., A. Maki, and M. A. Mazurowski (2018): "A systematic study of the class imbalance problem in convolutional neural networks," *Neural Networks*, 106, 249 – 259.

Carroll, C. D., T. F. Crossley, and J. Sabelhaus (2015): *Improving the Measurement of Consumer Expenditures*, University of Chicago Press.

Dahl, G. B., K. V. Løken, and M. Mogstad (2014): "Peer Effects in Program Participation," *American Economic Review*, 104, 2049–2074.

De Giorgi, G., A. Frederiksen, and L. Pistaferri (2016): "Consumption Network Effects," Working Paper 22357, National Bureau of Economic Research.

Deng, J., W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei (2009): "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*.

Gebru, T., J. Krause, Y. Wang, D. Chen, J. Deng, E. L. Aiden, and L. Fei-Fei (2017): "Using Deep Learning and Google Street View to Estimate the Demographic Makeup of the US," *arXiv:1702.06683 [cs]*, arXiv: 1702.06683.

Gilchrist, D. S. and E. G. Sands (2016): "Something to Talk About: Social Spillovers in Movie Consumption," *Journal of Political Economy*, 124, 1339–1382.

Grinblatt, M., M. Keloharju, and S. Ikäheimo (2008): "Social Influence and Consumption: Evidence from the Automobile Purchases of Neighbors," *The Review of Economics and Statistics*, 90, 735–753.

Haenssle, H. A., C. Fink, R. Schneiderbauer, and [et al.] (2018): "Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists," *Annals of Oncology*, 29, 1836–1842.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102, 994–1028.

Krizhevsky, A., I. Sutskever, and G. E. Hinton (2012): "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems 25*, ed. by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Curran Associates, Inc., 1097–1105.

Kuhn, P., P. Kooreman, A. Soetevent, and A. Kapteyn (2011): "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery," *American Economic Review*, 101, 2226–2247.

Lavy, V., O. Silva, and F. Weinhardt (2012): "The Good, the Bad, and the Average: Evidence on Ability Peer Effects in Schools," *Journal of Labor Economics*, 30, 367–414.

Manski, C. F. (1993): "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, 60, 531–542.

Moretti, E. (2011): "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales," *Review of Economic Studies*, 78, 356–393.

Szegedy, C., S. Ioffe, and V. Vanhoucke (2016): "Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning," *CoRR*, abs/1602.07261.

Tibshirani, R. (1996): "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, 58, 267–288.

Vasconcelos, C. N. and B. N. Vasconcelos (2017): "Increasing Deep Learning Melanoma Classification by Classical And Expert Knowledge Based Image Transforms," *CoRR*, abs/1702.07025.

A Appendix

A.1 Image prepossessing

In this section, I detail the data collection process. To obtain the trampoline ownership status of the homes in my sample using a neural network, a number of preprocessing steps is necessary. The aerial photographs used are produced by the Swedish mapping, cadastral and land registration authority (Lantmäteriet). The set of images I use covers the major population centers and the surrounding areas (table 7 presents the distribution of land lot-year observations at the municipal level). The photographs are taken during 2006 until 2017. The observation weighted mean time between consecutive photos of a land lot is 2.5 years. A land lot in my sample is on average photographed 4.2 times. The images are orthorectified, meaning that the photo has been projected on an elevation model of the earth's surface and a standard mapping coordinate system is embedded in the image. This allows maps and other spatial data to be accurately projected on the image. The images are typically taken during spring, just after trees and bushes start to leaf out. The median photo date for 2010 until 2017 is May 19$^{\text{th}}$.

To detect the presence of a trampoline on a specific lot, the first prepossessing step is to cut the large [20] aerial photos down to small image chips, each containing a single land lot. As an input, the neural net accepts a fixed three-dimensional matrix. In this study, I set the shape of the input as (300,300,3), which corresponds to a 300 by 300 pixel using the standard 24-bit three-channel RGB color space[21]. The aerial photos have a spatial resolution of 0.25 meters per image pixel, meaning that the largest eligible dimension for a land lot is 75 by 75 meters ($300 * 0.25$). This is sufficiently large to fit the majority of single-family home lots and only about 2 percent of eligible lots have to be discarded for being too large. Increasing the input size increases the computational burden exponentially and scaling down images of larger lots to fit the input size would decrease the pixel size of trampolines, making them more difficult to detect.

It is crucial that each image only shows a single land lot, or we risk contaminating the input with e.g. neighboring trampolines. To achieve this, I use a vector based map of land lot borders produced and maintained by Lantmäteriet. Projecting this map on the aerial photo allows me to cut along the borders to isolate a single property. As I am only interested in properties occupied by single-family homes, I overlay a map of building footprints and their legal building code and only include land lots that intersect a build registered as a single-family home. While land lots

[20] The original images are 10,000 × 10,000 pixels.

[21] The third dimension corresponds to the red, green and blue color space. Each 300 by 300 channel takes an integer value between 0 and 255 (8 bits).

with multiple family homes are rare, I have to exclude these since I would not be able to attribute a trampoline to a single household. Figure 9 illustrates the property selection procedure. Since the image is always cut along the x-y axis of the original photo (roughly corresponding to the east-west and north-south axis of the map), the cardinal alignment of the land lot matters for the dimensions of the cutout. To avoid selecting my sample on this rather arbitrary factor, I implement a simple rotation algorithm. Before discarding a land lot as being too large, the algorithm attempts to rotate lots that exceed the 300 pixel restriction in either dimension up to 90 degrees. Figure 11 illustrates the rotation procedure.

Because the map detailing property borders is continuously updated and only the latest version is made available, there are discrepancies between the ground truth at the time the photo was taken and the border map. For example, a property observed in 2017 may not have been formed and built upon back in 2006. Consequently, the lot cutout from the 2006 photograph would just depict an empty plot of land. Since a lot has be occupied in order to meaningfully be classified into having or not having a trampoline, this requires me to first classify lots as occupied or empty. I address this ex ante classification problem using a secondary neural net, trained on images containing occupied and empty lots. I manually label a set of 9,618 images (3,084 empty and 6,534 occupied lots). The set of labeled images is balanced and split 85/15 into a training and validation data set. After training, the neural network achieves 97.6 percent accuracy in the validation data set. The trained network is then used to predict the full set of images. If the predicted probability of being occupied is 0.5 or above, I consider the land lot to be occupied in the current and all future years. Because new housing is marginal compared to the existing stock, the vast majority of lot-year observations (98 percent) are classified as occupied. As changes to this secondary neural network only concerns sample selection along a very small margin, I do not detail the performance of this network further.

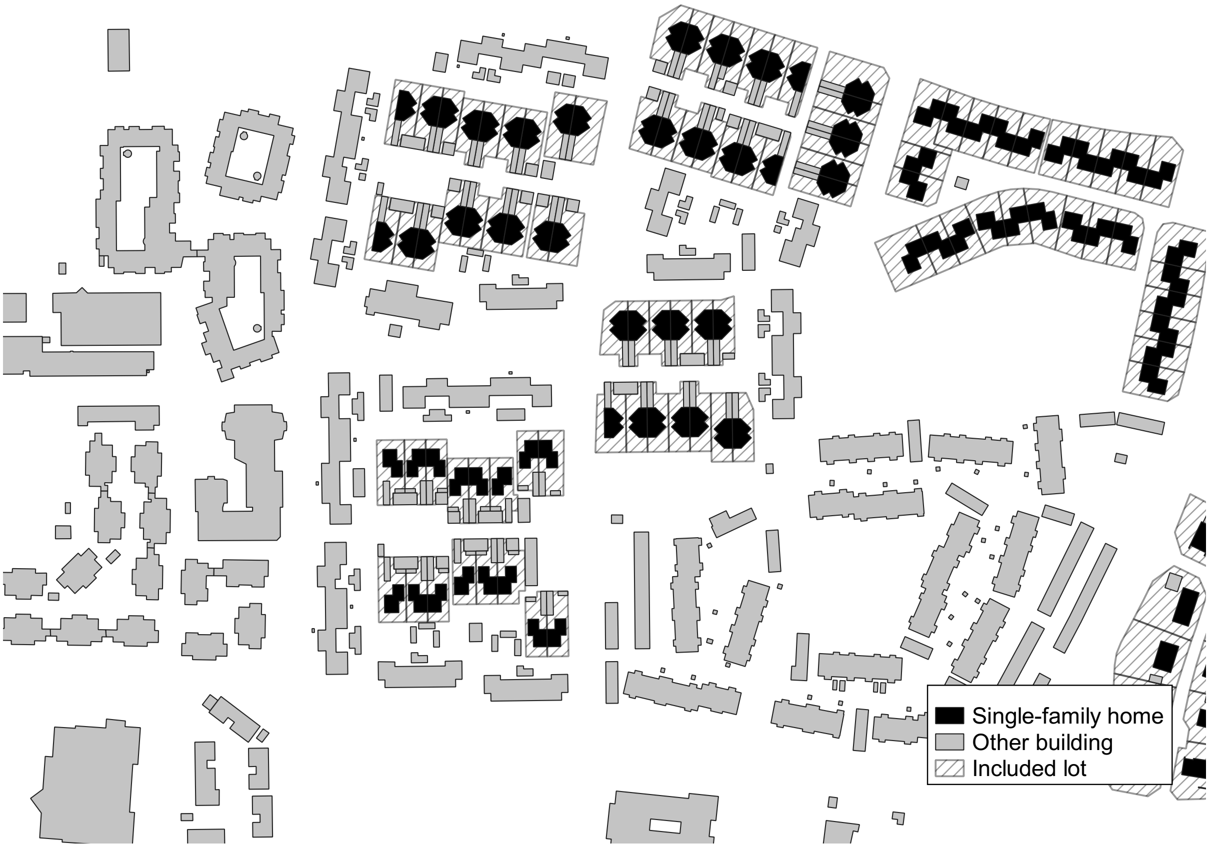

Figure 9: Black shapes represent the footprint of single family homes. Their corresponding land lot is included in my sample (striped areas).

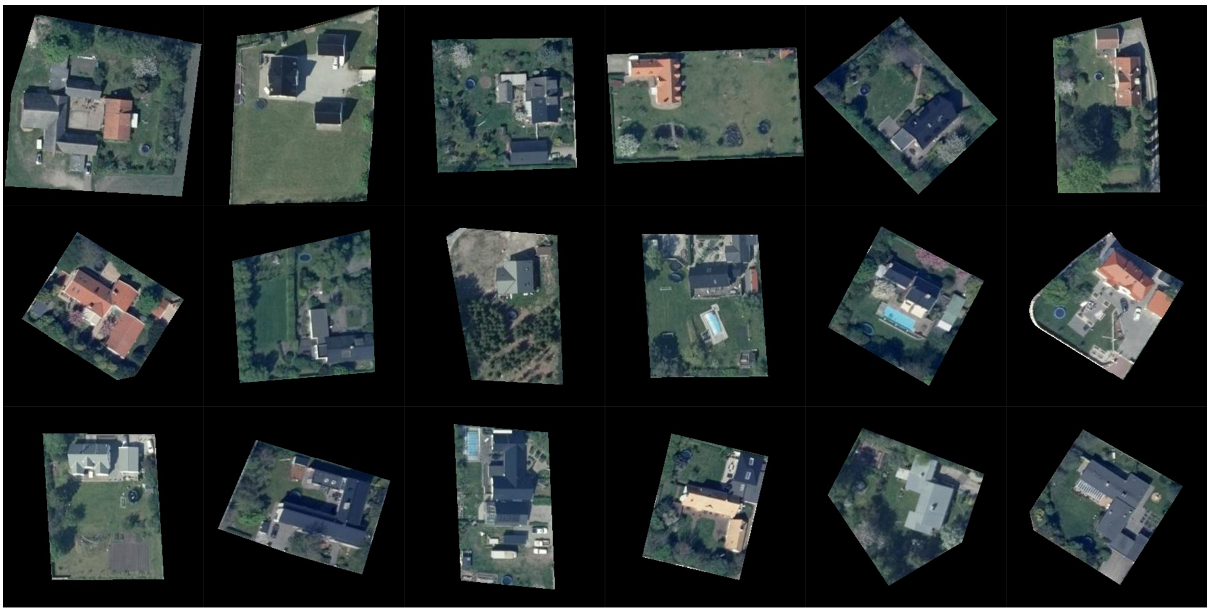

Figure 10: Examples of images fed into the neural network. Each property is a separate 300 by 300 pixel image, corresponding to 75*75 meters.

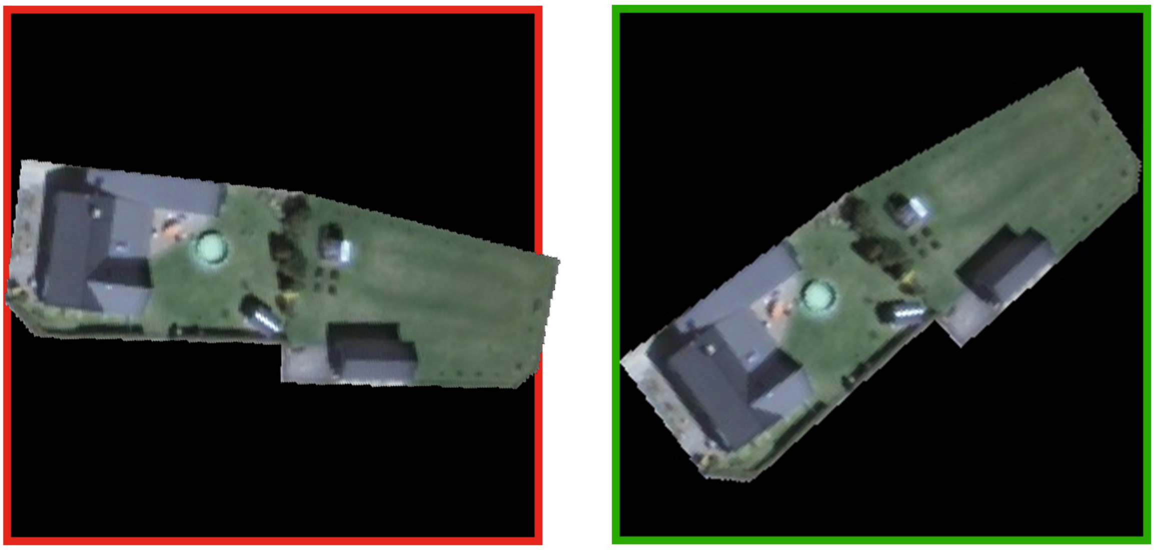

Figure 11: An example of an image violating the 300 pixel restriction due to the x-y alignment of the land lot (left). After a 43 degree counter-clockwise rotation, the lot satisfies the restriction (right).

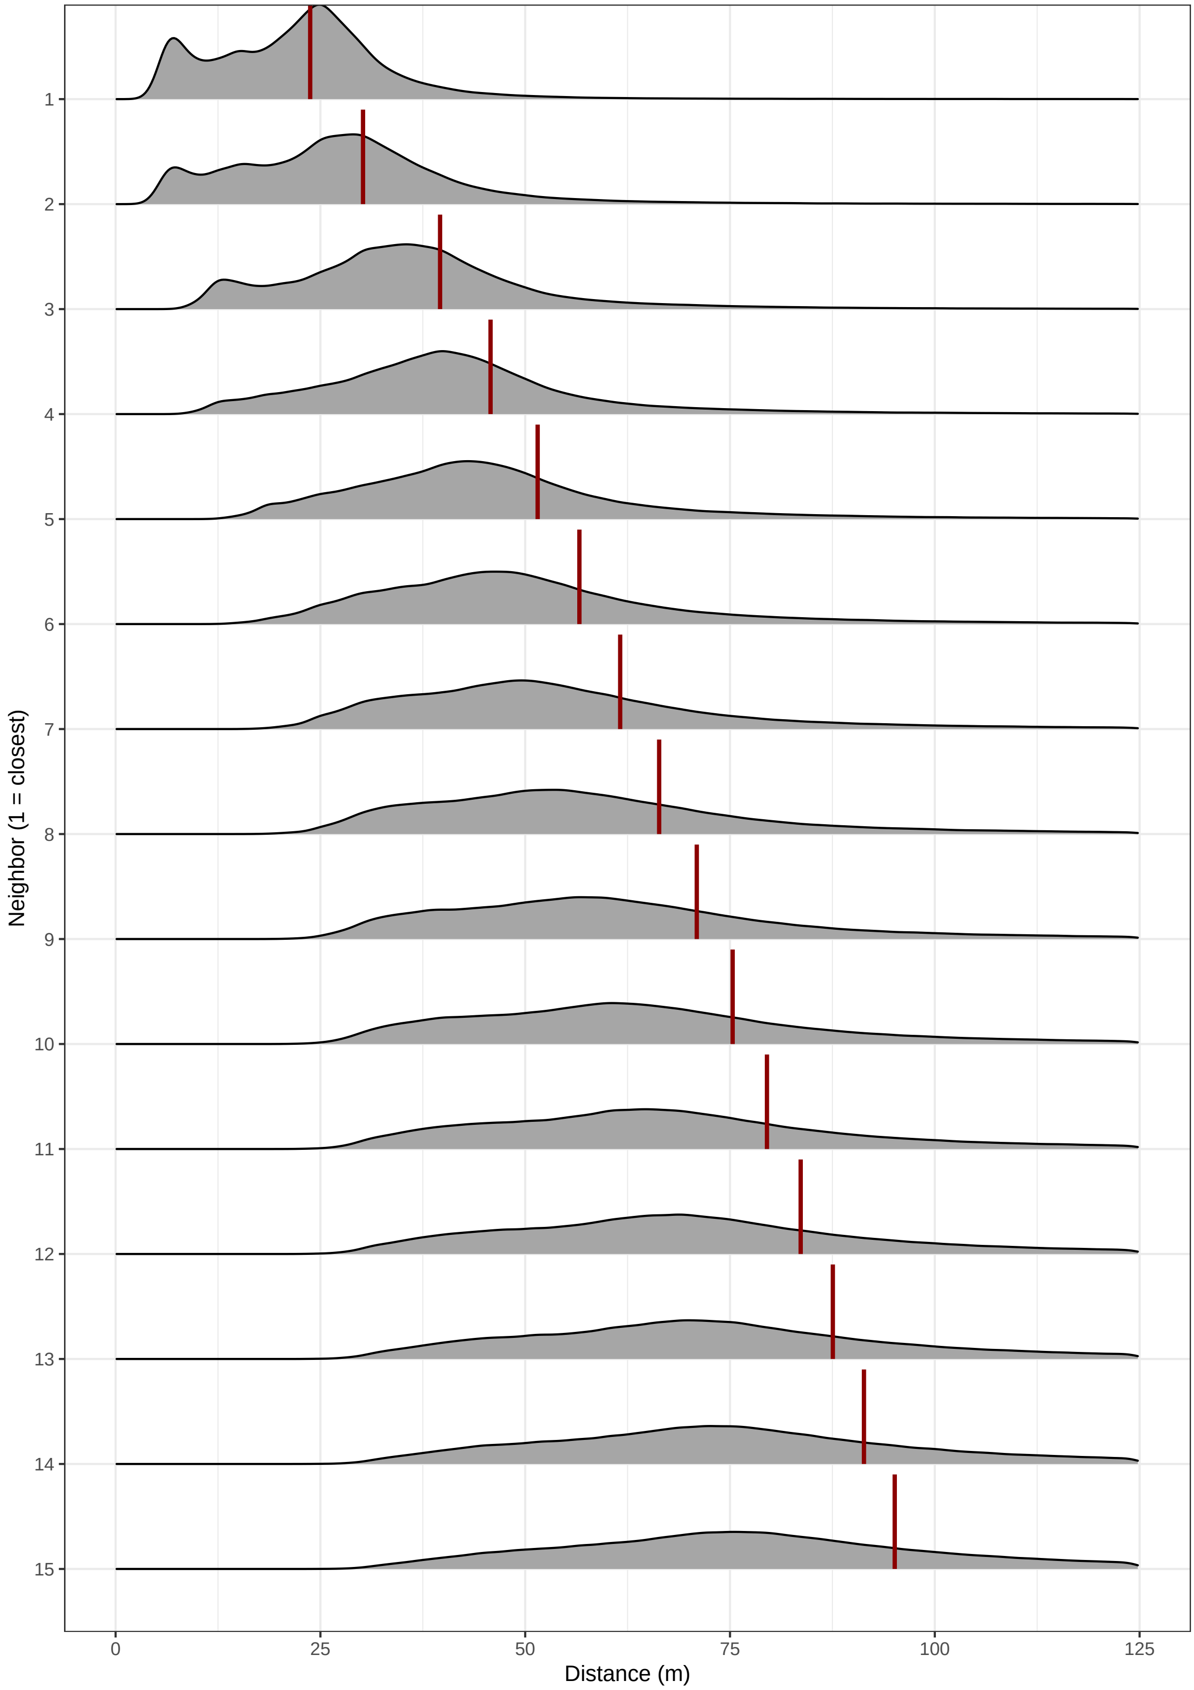

Figure 12: Density plots of the sample centroid-to-centroid distance to the 15 closest neighbors. The red lines indicate means.

Table 7: Observations by municipality

Municipality	No.	%	Municipality	No.	%
Ale	4,644	0.22%	Mark	8,796	0.41%
Alingsås	8,632	0.40%	Mjölby	1,559	0.07%
Bjuv	8,941	0.42%	Motala	1,505	0.07%
Borås	37,663	1.75%	Mullsjö	2,944	0.14%
Botkyrka	38,335	1.78%	Mölndal	38,174	1.78%
Burlöv	14,542	0.68%	Nacka	53,223	2.48%
Danderyd	24,829	1.16%	Norrköping	45,807	2.13%
Ekerö	15,694	0.73%	Nykvarn	4,829	0.22%
Eskilstuna	37,996	1.77%	Nynäshamn	1,065	0.05%
Eslöv	902	0.04%	Nässjö	1,705	0.08%
Finspång	8,209	0.38%	Partille	25,130	1.17%
Forshaga	1,722	0.08%	Salem	11,455	0.53%
Gnesta	2,785	0.13%	Sigtuna	14,845	0.69%
Göteborg	183,445	8.54%	Sjöbo	1,021	0.05%
Habo	3,530	0.16%	Skurup	5,237	0.24%
Hallsberg	547	0.03%	Sollentuna	42,276	1.97%
Hallstahammar	8,269	0.38%	Solna	1,432	0.07%
Hammarö	14,999	0.70%	Staffanstorp	19,957	0.93%
Haninge	43,996	2.05%	Stockholm	158,402	7.37%
Helsingborg	68,849	3.20%	Strängnäs	3,899	0.18%
Huddinge	65,306	3.04%	Sundbyberg	4,889	0.23%
Härryda	23,816	1.11%	Surahammar	3,153	0.15%
Håbo	14,306	0.67%	Svalöv	4,282	0.20%
Höganäs	5,465	0.25%	Svedala	12,187	0.57%
Höör	9	0.00%	Söderköping	1,886	0.09%
Järfälla	36,453	1.70%	Södertälje	28,412	1.32%
Jönköping	66,700	3.10%	Trelleborg	3	0.00%
Karlstad	42,254	1.97%	Tyresö	31,134	1.45%
Kil	3,723	0.17%	Täby	53,380	2.48%
Klippan	8,316	0.39%	Umeå	47,363	2.20%
Knivsta	5,240	0.24%	Upplands Väsby	20,518	0.96%
Kristinehamn	12,356	0.58%	Upplands-Bro	12,235	0.57%
Kumla	12,242	0.57%	Uppsala	54,029	2.51%
Kungsbacka	12,090	0.56%	Vaggeryd	1,849	0.09%
Kungsör	2,497	0.12%	Vallentuna	19,354	0.90%
Kungälv	14,454	0.67%	Vaxholm	8,627	0.40%
Kävlinge	16,234	0.76%	Vellinge	7,138	0.33%
Köping	8,626	0.40%	Vännäs	1,624	0.08%
Landskrona	19,737	0.92%	Värmdö	21,429	1.00%
Lekeberg	993	0.05%	Västerås	63,216	2.94%
Lerum	30,458	1.42%	Ängelholm	95	0.00%
Lidingö	26,382	1.23%	Åstorp	11,200	0.52%
Linköping	55,359	2.58%	Öckerö	7,329	0.34%
Lomma	23,217	1.08%	Örebro	40,388	1.88%
Lund	47,675	2.22%	Österåker	33,003	1.54%
Malmö	134,031	6.24%	Total	2,148,452	100.00%