Overweighting Product Quality Information from a Single Unrepresentative Peer: Evidence from Hyper-Local Neighborhoods

Sumit Agarwal, Souphala Chomsisengphet, Sahil Raina, Barry Scholnick, and Man Zhang*

March 19, 2025

Abstract

We show that individuals overweight product quality information from a single peer, despite this information being unrepresentative. Using geocoded data on purchases and warranty claims of consumer durable goods, and a hyper-local spatial difference-in-differences identification strategy, we show that individuals reduce consumption of a brand when a single neighbor makes a warranty claim on that brand. A single warranty claim causes a 7% shift to other brands by close neighbors of the claimant. This evidence shows the importance of single unrepresentative peers in consumption decisions.

Keywords: Consumption Peer Effects, Product Quality
JEL Codes: D83, D12, G52

*Agarwal: National University of Singapore, ushakri@yahoo.com; Chomsisengphet: Office of the Comptroller of the Currency, US Treasury, Souphala.Chomsisengphet@occ.treas.gov; Raina: Alberta School of Business, University of Alberta: sraina@ualberta.ca; Scholnick: Alberta School of Business, University of Alberta: barry.scholnick@ualberta.ca; Zhang: Institute of Real Estate and Urban Studies, National University of Singapore: man.zh@nus.edu.sg.The views expressed in this paper do not represent those of the Office of the Comptroller of the Currency, or the U.S. Department of the Treasury. Raina and Scholnick thank the Social Sciences and Humanities Research Council of Canada (SSHRC) for funding. We thank participants at the European Urban Economics Association, the Consumer Finance Decision Making Conference, the Urban Economics Association, and the Federal Reserve Bank of Philadelphia

1 Introduction

Individuals should not pay attention to signals arising from a single peer, particularly if the peer's information is based on unrepresentative experiences. Akerlof in his 1991 Ely Lecture, however, argues that individuals may disregard this basic insight and "attach too much weight" (p. 2) to the negative consumption experiences of a single salient peer when making their own consumption decisions, even though this information from the single peer may be idiosyncratic, and not representative of the true quality of a good.

In his lecture, Akerlof presents a hypothetical anecdote in which a consumer searching for a car determines that Volvo is a high quality brand. However, a subsequent interaction with one of her friends, who shares a recent negative experience with the quality of Volvo cars, pushes the consumer away from purchasing a Volvo, even though the information from the friend is likely biased due to the friend's recent idiosyncratic and unrepresentative experience with the brand. Akerlof describes this anecdote as reflecting a "central principle of modern cognitive psychology" (p. 2). This "cognitive psychology" view is in opposition to the generally accepted belief that individuals should *ignore* the information shared by a peer whose experience and opinion is idiosyncratic and likely unrepresentative of the truth. The aim of this paper is to test the validity of the "cognitive psychology" view by assessing whether individuals pay attention to unrepresentative opinions of their peers.

Specifically, our paper examines the effect on an individual's consumption decisions of a single salient peer's warranty claim on a consumer durable good (e.g., refrigerators, dishwashers, laundry machines, etc.). We examine how such a warranty claim by a single peer (the sending peer) impacts the future purchases of the claimed brand's products by other peers (i.e. the receiving peers). Our main finding is that a warranty claim filing by the sending peer, significantly lowers subsequent purchases of the claimed brand's products by the receiving peers.

We argue that there are two main reasons for why our use of a single warranty claim as an information shock is appropriate for testing the "cognitive psychology" view of Akerlof above. First, the fact that the information is transferred from a single sending peer implies that the information will be *idiosyncratic* (i.e. based on a sample size of one). Second, we argue that the receiving peer should be aware that information from a warranty claim will not, on average, be *representative* of the true quality of the product. This is because the provision of warranties by producers is typically taken to signal the high quality of the product, i.e., the provision of warranties by producers can be

taken to reflect the costly signaling of quality by the producer (Dranove and Jin, 2010). Indeed, our data show that warranty claims are very rare, with approximately 1.5% of all consumer durable sales resulting in a warranty claim. Thus in order for the receiving peer to change consumption choices based on the information shock from a single warranty claim, would require the receiving peer to ignore both that this information is idiosyncratic and, importantly, unrepresentative of product quality for that product.

These features of our setting make peer responses to warranty claim filings an ideal setting to examine the "cognitive psychology" hypothesis of Akerlof. Our main finding in this paper, that receiving peers do indeed reduce consumption of a brand, following a warranty claim on that brand by a single sending peer, is thus consistent with the "cognitive psychology" proposition that consumers do overweight such idiosyncratic and unrepresentative information.

Our data are provided by a major Canadian consumer durable goods retail store. This retailer provided us with detailed information on all purchases of and warranty claims filed on consumer durable goods between 1998 and 2009, which amounted to six million sales transactions across over 200 brands along with 95,000 warranty claim filings. Each observation includes details such as purchase and warranty claim dates and purchaser and claimant home addresses.

To identify the causal effects of warranty claimants' negative experiences on peers' purchasing behavior, we define claimants' peers based on geographic proximity. We leverage the precise geographic location of consumers and precise claim and purchase dates provided by our data to conduct an empirical analysis employing a hyper-local difference-in-differences empirical methodology. In particular, we compare the effect of a claim on the claimant's immediate neighbors' purchasing behavior to its effect on the purchasing behavior of slightly more distant neighbors. This identification strategy has been used extensively in many areas of economics, where treatment is argued to decline with distance. This large literature, described in detail below, exploits the richness of fine-grained geocoded spatial data at hyper-local distances. In our specification, we examine how information flows from the single information-sending peer (the warranty claimant) to information-receiving peers who are close neighbors of the information sender.

The main identification assumption of this hyper-local identification strategy is that very close (treatment group) neighbors of the information sender are much more likely to receive the information from the sender and (control group) neighbors who are "slightly further away" from the information sender are less likely to receive the information. Following the hyper-local identification literature, we further assume that whether a household chooses to live in the treated group or

in the control group is as good as random. This assumption is reasonable in our setting as housing markets are typically thin at granular scales when a household originally arrives in a neighborhood. By this second assumption, given the very close distances between the treatment and control group households, it is reasonable that treatment and control group neighbors have similar observable and unobservable characteristics and are exposed to the same potentially confounding events that may influence purchasing decisions. Therefore, this identification strategy can address endogeneity issues such as omitted variables and unobserved common shocks.

As in all peer effects studies, we need to address the well-known "reflection problem" of Manski (1993), which essentially concerns pinning down the exact direction of information flow between peers. To address the reflection problem, we adopt a strict sample selection strategy, whereby our analysis sample only includes neighborhoods with: (1) exactly one warranty claim event over the entire sample period (i.e., exactly one information-sending neighbor), and (2) no households that fall into the neighborhoods of any other warranty claimant. These sample selection procedures effectively require that any information-receiving household (in either the inner or the outer ring) can only receive information about the warranty claim from exactly one information-sending close neighbor (the warranty claimant). This sample selection criteria, restricting our sample to neighborhoods with exactly one warranty claim event, thus addresses the reflection problem and allows us to cleanly test the "cognitive psychology" hypothesis posed in Akerlof (1991), that individuals respond to (unrepresentative) information from a single salient peer.

Our main baseline finding is that a warranty claim filing significantly reduces the consumption of the claimed brand's products relative to other brands' products by inner ring (treatment group) households relative to outer ring (control group) households. In terms of magnitudes, we find that, in our preferred baseline specification, in the four quarters after a warranty claim, relative to the four quarters before the claim, inner ring neighbors (residing in postal codes from 0 m to 100 m away from the claimant's postal code) make 1.6 fewer net purchases of the claimed brand, when compared to outer ring neighbors (residing from 100 m to 200 m away from the claimant). This amounts to about 7% of consumer durable goods being purchased from a different manufacturer due to the warranty claim, an economically meaningful shift in the consumption behavior of a claimant's peers.

There is a potential concern that the effect we attribute to idiosyncratic, unrepresentative information shared by the warranty claimant is instead driven by some component of the warranty claimant's signal that is correlated with true measures of product quality. To examine this alterna-

tive hypothesis, we create three measures of objective product quality, aggregate claims rate, brand market share, and warranty repair time, and include them as explanatory variables in our regression analysis. If information correlated with these objective quality measures is the real driver of neighbors' purchasing behavior, then including them in our regressions should significantly weaken the impact of the warranty claimant's information. When we run this analysis, we find that effect of the warranty claim on neighbors' purchasing behavior is unchanged by the inclusion of controls. This finding confirms that the warranty claimant's effect on neighbors' purchases is driven by their own idiosyncratic, unrepresentative experience, as proposed in the "cognitive psychology" view posed by Akerlof.

In many of the papers that study peer effects, the impact of the peer effect is shown to be mediated by the strength of the social network among those peers. As such, for our geography-based peer effect, it is possible that neighborhood characteristics that affect the strength of the social network among neighbors intermediates the effect of the warranty claim on the claimant's neighbors' purchasing behavior. We examine this possibility for three neighborhood characteristics that may have an effect on the strength of the neighborhood's social network: average income, population density, and immigrant proportion. We find that all three neighborhood characteristics influence the magnitude of the claim's peer effect, which implies that the neighborhood social network intermediates the peer effect of the warranty claim on neighbors' purchasing decisions.

Our study contributes to various literatures. A large literature has examined consumption peer effects in a variety of contexts.[1] However, as described in detail above, a study of the "cognitive psychology" view of Akerlof (1991) requires two main characteristics: (1) that the focus be on the effect of a single sending peer; and (2) that the single sending peer send unrepresentative information to any receiving peer. To the best of our knowledge, no other consumption peer effects paper studies such a setting. While the majority of the studies cited here examine multiple sending peers, some of studies examine settings with a single sending peer.[2]

[1]Studies examining peer effects in consumer products include Agarwal et al. (2020, 2021); D'Acunto et al. (2024); De Giorgi et al. (2020); Bertrand and Morse (2016); Bailey et al. (2022); Moretti (2011); Gilchrist and Sands (2016) and financial products (Brown et al., 2008; Bursztyn et al., 2014; Duflo and Saez, 2003; Heimer, 2016; Hong et al., 2004, 2005; Ouimet and Tate, 2020; Pool et al., 2015; Maturana and Nickerson, 2019; Hacamo and Kleiner, 2022).

[2]Examples of consumption peer effects studies that focus on a single sending peer include studies of income shocks from single lottery winners (Agarwal et al., 2020; Kuhn et al., 2011), income shocks from single bankruptcy filers (Agarwal et al., 2021) and RCT studies (Bursztyn et al., 2014). Our study is also similar to Bailey et al. (2022), in that both examine how a single peer sends product quality information to a receiving peer.

In terms of the transmission of information that is unrepresentative of the aggregate, we argue that our setting is very different from the existing consumption peer effects literature. The existing consumption peer effects literature almost universally examines how *positive* information sent by the sending peer causes the receiving peers to increase their consumption of that product. In our warranty claim setting, on the other hand, we examine a negative information event, reflecting low product quality, causes a reduction in consumption by the peer. To extend the metaphor of Akerlof (1970), we focus on negative product information transfer between peers (i.e., where the products are "lemons") while most of the existing literature examines information transfer between peers regarding positive product information (i.e., where the products are "peaches"). We argue that a positive experience of a peer is more likely to be an accurate reflection of the typical experience for that brand and, thus, it is typically not possible to make the claim that resultant "peach" peer interactions are based on unrepresentative information. On the other hand, in our specific warranty claim setting, we can make the plausible claim that the warranty claimant's experience is unrepresentative of aggregate quality. As we argued above, this is because consumers should understand that warranties are provided by producers as a costly signal, that the durable will not, on average, require a repair. [3]

This distinction between positive information causing increases in peer consumption (peaches) and negative information causing reductions in peer consumption (lemons) has recently been highlighted by Han et al. (2023). They argue that "there is evidence of contagion of consumption ..., but contagion can potentially spread either a decision to consume more or to consume less. Notably, little is known about whether biased social learning implies over- versus underconsumption" (p. 1648). Our paper is one of the few that provides causal evidence that peer effects (following a negative information shock) will cause a reduction in consumption by the receiving peer.

Our study also contributes to the literature on the measurement of product quality. While it is widely agreed that product quality is crucial for consumption decisions, a major challenge in empirical studies of consumption is that product quality is typically unobservable to the econometrician.[4]

[3]This is confirmed by the rarity of warranty claim filings, which, in our consumer durable good setting occur for only 1.5% of all durables that are purchased.

[4]As product quality is challenging to observe, various measures have been used in the literature, including; (1) self-reported quality by consumers (e.g., Busso and Galiani, 2019), (2) the evaluation of quality by a "skilled quality assessor" (e.g., Atkin et al., 2017), (3) the assumption that brands with a higher market share conditional on price have a higher quality (e.g., Khandelwal, 2010), and (4) experimentally varying quality in an RCT (e.g., Bold et al., 2022).

Our paper introduces a new measure of product quality to the literature: warranty claim filings. Unlike much of the existing literature, our warranty claim measure of product quality does not require subjective opinions (e.g., from consumers or from "expert" third parties). The one caveat with our new warranty claim measure of product quality is that an individual consumer's perception of quality could be multi-faceted, whereas we only observe one, albeit important, element of quality: the emergence of a post-purchase defect in a product that results in a warranty claim.

As argued by Kuchler and Stroebel (2021), consumption peer effects are not only important for understanding consumer choices, but also for understanding market-level outcomes for brands. One final important contribution of our paper is that we directly address the impact of peer effects on market-level outcomes because we measure the impact of the peer effect on claimed brand sales relative to all other brands. Bailey et al. (2022) is the only other study which examines how consumption peer effects generate changes in market-level outcomes.[5] As described above, an important difference between Bailey et al. (2022) and our paper is that they examine a "peaches" setting, while we examine a "lemons" setting.

2 Data

Our data consists of the universe of sales transactions as well as the universe of warranty claims made by purchasers for an anonymous Canadian retailer of consumer durable goods, spanning the period from 1998 to 2009.[6] The durable goods retailer has a market share that is among the top four in Canada. It has more than 270 individual stores (which are typically "big box" stores) across all regions in Canada. Prior to data cleaning procedures, our database contains 6.54 million sales transactions of more than 35,000 different products to over 3 million consumers.

2.1 Sales Data

The sales transactions data provide us detailed information for each product purchased from the retailer. The dataset includes attributes such as product brand, product category (e.g., kitchen, laundry, TV, video, stereo), product code (e.g., refrigerator, dryer), model series number, price, date of purchase, and a unique identifier for each sales transaction.

[5]Bailey et al. (2022) argue that "within the literature that has studied peer effects in product adoption decisions, we are the first, to our knowledge, to identify important competitive spillovers to other models and brands." (p. 492)

[6]This data is also used in Bošković et al. (2024).

During the data cleaning process, some observations were dropped due to missing information or data imputation errors. As a result, our cleaned dataset consists of 5.9 million observations on individual sales transactions. The summary statistics of the sales data are presented in Table 1a. Our sales data also includes a unique customer identifier, which allows us to track each consumer's purchases over time. The consumer has an incentive to provide identifying information (including name and address) at the point of sale in order to be able to claim on a warranty for that product.

While all products are typically covered by a manufacturer's warranty (of approximately one year), the consumer also has an option to purchase an extended warranty. As can be seen in Table 1a, an extended warranty is purchased in 36% of all sales transactions. Importantly, regardless of whether a warranty claim is made on a manufacturer's warranty or an extended warranty, the claim is treated identically (i.e., which products will be repaired, who conducts the repairs, and the extent of repairs). As such, in this paper we treat them as interchangeable.[7]

A key element of our sales data is our ability to observe the brand of every product sold. Our sales data includes more than 200 brands, with Table 1b providing a listing of the top 30 brands. Our ability to observe brands for each transaction is central to our methodology because, as described in detail below, our dependent variable consists of the number of sales of the claimed brand minus the number of sales of all other brands in each neighborhood.

2.2 Warranty Claims Data

Our dataset also includes all the warranty claims filed during the sample period. A warranty claim record can be linked with the sales transactions data as described above using the unique sales transaction identifier. Within the warranty claims data, we have access to the exact date when each warranty claim occurs, which we define as the claim event date. The summary statistics of all claim events are reported in Table 1c. After cleaning, we have approximately 95,000 claim events in our data.

As explained in detail below, for reasons related to the identification of peer effects and the reflection problem, we only retain a subset of claim events for our analyses.[8] Due to our strict

[7] Bošković et al. (2024) uses this data to examine the pricing of extended warranties, while in this paper we examine how a warranty claim (on either an extended or manufacturers warranty) impacts neighborhood peer effects.

[8] In particular, we only retain claim events where there is a single claim recorded in that neighborhood and no treated or control neighbors appear in neighborhoods of other warranty claims. We explain our sample selection criteria in further detail in Section 4.

sample selection criteria, we are left with 1,776 warranty claim events for our analyses. The warranty claims data summary statistics reported in Table 1c show that the full data sample of claim events and the analysis sample exhibit similar observable characteristics.

2.3 Geographic Neighborhood Data

For each consumer, we observe their home address's six-digit postal code. Canadian six-digit postal codes represent extremely small geographic areas, with the average postal code containing approximately 19 households, and the median postal code containing 13 households. We use the six-digit postal code to identify the geographic location of every individual in our study. Because we observe the postal code of every individual in our study, we can calculate the geographic distance between all individuals in our study. As we describe below, we use these geographic distances to define very close (inner ring) neighbors of warranty claimants as treated households and slightly further away (outer ring) neighbors as control households. A graphical description of these six-digit postal codes, along with our preferred ring sizes (inner ring treated ring from 0 m to 100 m, and outer control ring from 100 m to 200 m) is provided in Figure 1. A number of other papers (e.g., Agarwal et al., 2020; Baum-Snow et al., 2024) also use Canadian six-digit postal code data to measure the locations of individuals as well as the distances between individuals.

In addition to the six-digit postal code, we also make use of slightly larger geographic areas known as a Dissemination Areas (DAs). Each DA has a population of between 400 and 700 individuals. DAs are the smallest geography in Canada for which we observe detailed census data. We match every six-digit postal code to the appropriate DA using the Post Code Conversion File provided by Statistics Canada to obtain demographics for our sample.

The summary statistics of the DA-level demographic data are shown in Table 1d. In this table, we compare the demographic variables for DAs that are included in our main analysis sample (based on the sample selection criteria described below) with all other DAs in Canada. As can be seen in this table, the observable demographic characteristics for DAs included in our analysis are very similar to those demographic characteristics of all the other DAs in Canada.

3 Causal Identification

3.1 Ring Method for Hyper-local Differencing

In our setting, we are interested in the flow of information (about product quality) via word-of-mouth between a single sending peer (the warranty claimant) and the receiving peers (her close neighbors). Given the spatial nature of our question, we follow the hyper-local differencing procedure for defining treatment and control group neighbors. This approach has become standard in empirical studies where treatment is a function of distance; i.e., in settings where units close to the focal event can be considered as treated by the event, while units slightly further away from the event are not treated, and can be used as controls. This spatial approach to identification is sometimes referred to as the "ring method," because rings are drawn based on distance from the focal event to differentiate between treated units (who are in the inner ring, close to the event) and control units (who are in the outer ring, further away from the event). A graphical illustration of this hyper local ring method approach to identification in our setting, is provided in Figure 1.

A substantial literature in various branches of economics has used this ring-based spatial identification procedure, in contexts such as: (1) real estate (e.g., where distance to events such as neighborhood foreclosure and the location of sex offenders, impacts house prices)[9], (2) health (e.g., where the distance to local pollution sources impacts health outcomes)[10], (3) crime (e.g., where distance to police shootings impacts education outcomes) [11], productivity spillovers (e.g., where distance to other firms impacts firm level productivity)[12], firm competition (e.g., where distance to other firm impacts competition)[13], transit usage (e.g., where distance to transit impacts transit usage)[14] and neighborhood sorting (e.g., where distance to neighbors of other racial characteristics or political preferences causes home owners to exit neighborhoods).[15] Importantly, this approach has also been used in previous peer effects studies, where, as in our setting, the peers are neighbors

[9] Anenberg and Kung (2014); Bayer et al. (2021); Gupta (2019); Pool et al. (2015); Campbell et al. (2011); Linden and Rockoff (2008)

[10] Currie et al. (2010, 2011); Currie and Tekin (2015); Currie et al. (2015)

[11] Ang (2020)

[12] Rosenthal and Strange (2003); Arzaghi and Henderson (2008); Greenstone et al. (2010); Baum-Snow et al. (2024)

[13] Ellickson and Grieco (2013); Busso and Galiani (2019); Arcidiacono et al. (2020); Schiff et al. (2023)

[14] Gupta et al. (2022); Jerch et al. (2024)

[15] McCartney et al. (2024); Bayer et al. (2022)

[16].

Our study, like all of the aforementioned literature, exploits fine-grained geographic location data to assign households to treatment and control groups. We implement this neighbor assignment process using the fine-grained six-digit postal codes available to us. Our Canadian neighborhood setting is similar to papers such as Agarwal et al. (2020); Baum-Snow et al. (2024), which also use Canadian six-digit postal codes when implementing a hyper-local identification strategy, though those papers study questions quite different from ours.

3.2 Warranty Claims as Information Shocks

One of the main threats to identification of consumption peer effects comes from homophily within peer groups that have been endogenously formed. Because the peers have endogenously selected into their peer group, any observed similarities between peers (e.g., both peers consume the same product) could be because of peer homophily rather than actual peer effects (i.e., where the sending peer, through some action, directly causes the receiving peer to alter her consumption). One way to overcome this threat to identification is to exploit a shock which impacts only the sending peer and not the receiving peer. The key identification assumption of this strategy is that the receiving peer can only learn about the shock through her interactions with the sending peer.

In our paper, we use the warranty claim filed by the sending peer as evidence of a negative shock experienced by that peer alone. The negative product quality information generated by this shock can only be transmitted to the receiving peer by their interactions with the sending peer. While the warranty claim event is endogenously determined for the sending peer, the warranty claim event is a plausibly exogenous information event for all receiving peers, whose consumption choices are the focus of our study. Our identification strategy thus sets up an environment where the receiving peer can only obtain product quality information based on the warranty claim event through contact with the sending peer.

3.3 Treatment and Control Group Neighbors

An important feature of all studies using this hyper-local methodology is the use of fine-grained geographic data to define treatment and control groups. This allows us to address two important econometric issues: (1) endogenous selection into treatment and control groups and (2) differences

[16] Agarwal et al. (2020); Bayer et al. (2021); Baum-Snow et al. (2024)

in neighborhood characteristics between treatment and control groups. We discuss each in turn.

3.3.1 Endogenous Selection into Treatment and Control Groups

An important challenge in defining treatment and control groups is to ensure that there is no systematic selection bias in the definition, i.e., that individuals do not (or cannot) choose to be a part of either the treatment or control group. This issue can be problematic in spatial or geographic empirical settings, where it is possible for an individual to endogenously select into living in a specific neighborhood. For instance, a household may choose to be in a specific country, region, city, etc., because the observable or unobservable characteristics of the people in that area (e.g., ethnicity, political beliefs) match those of the household.

Fortunately, our hyper-local differencing strategy is able to address this issue. As has been frequently argued by the literature, while the individual may endogenously select into a wider neighborhood because of the characteristics of the neighborhood, whether the individual chooses to live in specific sub-section of the neighborhood (such as the postal codes used to define our treatment and control groups) can be considered "as good as random." In our study, the distance between any two treatment and control group households for a given warranty claim is quite small (in our primary specification, the centroids of postal codes assigned to the treatment and control groups have to be in contiguous 100 m concentric rings around the claimant postal code). Often, control group households are on "the other side of the road" from their treatment group counterparts (as postal codes often change from one side of a road to the other). As such, while an individual may choose the wider neighborhood they live in, residential real estate markets are typically sufficiently thin at the scales at which our treatment and control groups are defined that the specific locations of the housing available within the neighborhood at the time that they were searching are nearly impossible to control. Therefore, an individual cannot reasonably select into our treatment or control groups and, as a result, their group assignment can be considered "as good as random."

3.3.2 Neighborhood Unobservable Characteristics

A related problem in any spatially-defined treatment and control groups is that unobservable characteristics of neighborhoods (e.g., a mass layoff by a large neighborhood employer or neighborhood-specific preferences for certain products or brands) could impact treatment and control groups differently and, if correlated with the variables of interest, threaten causal identification. Once again, our hyper-local methodology addresses this issue, precisely because of fine-grained spatial definitions of

treatment and control groups. Because of the smallness of the geographic areas defining treatment and control groups for each warranty claim, any unobservable neighborhood feature should impact both the treatment and control groups for that claim very similarly. For instance, a mass layoff, a gas leak, or neighborhood-specific preferences should exert a nearly identical influence on both the treatment and control groups, given their geographic proximity. Thus, we effectively nullify the impact of such unobservables by comparing treatment group households to very geographically proximate control group households.

3.4 Determining the Width of Treatment and Control Rings

A central challenge in this rings-based hyper-local methodology is that the true treatment status of neighbors is unobservable. We observe the location of each neighbor, rather than whether a neighbor interacts with the warranty claim-filing neighbor. As a result, it is unclear, a priori, where to draw the rings containing treatment and control groups. As discussed in Butts (2023), if the treatment ring is defined as an area that is "too narrow," then it is more likely to exclude units that are actually treated (i.e., neighbors that interact with the warranty claimant). Those treated neighbors will incorrectly be included in the outer control ring. On the other hand, if the treatment ring is defined as an area that is "too wide," then the treated ring will incorrectly include some neighbors who are not treated (i.e., neighbors that do not interact with the warranty claimant).

To address this issue, we closely follow the methodology suggested in Butts (2023). The paper proposes a non-parametric, data-based methodology which identifies the best line to separate treatment and control groups. Effectively, the methodology proposes creating a large number of concentric rings of increasing distance and examining if and how the estimated treatment effect declines with distance from the focal event. A graphical description of the specifics of this method is provided in Figure 2. This method uses a DID framework to compare the effect of the information event in every ring to a "last" (i.e., outermost) ring. This last ring is assumed to be sufficiently distant that the (unobservable) distance at which the distance-based effect of the information event disappears falls somewhere between the site of the focal event and the last ring. In Figure 2, we illustrate this by showing 9 red inner concentric rings (labeled 1 to 9 with increasing distance from the claim), and then a last blue ring (labeled C).

The Butts method examines the DID coefficient estimate for each of the red inner rings in turn as the treatment group, using the last blue ring as the control group in all the DIDs. The key insight of the Butts methodology is that, using this series of DIDs where the treatment group

gets increasingly distanced from the information event, we can observe where the distance-based treatment effect (i.e. the DID coefficient) converges to zero. This point, where the series of DID coefficients for ever more distant red rings (all compared to the same last blue ring) converges to zero, can be considered a data-based estimate of the distance at which data-based effects of the focal event disappear. We can thus use this distance to define the treatment and control rings for our subsequent analyses.[17]

As we describe in detail in Section 5.3 below, we run the Butts (2023) procedure for our data. We find a very significant DID coefficient in the first ring, but that the estimated DID coefficients converge to zero and remain at zero for all subsequent rings. Based on evidence from this exercise, in all our main tests, we use households situated 0 m to 100 m away from the warranty claimant as our treated ring and households situated 100 m to 200 m away as our control ring.

4 Sample Selection

4.1 The Reflection Problem

The exogeneity of the information shock (the warranty claim) to receiving peers in both the treated and control groups is important for causal interpretation of our results. The well-known reflection problem first posed in Manski (1993) is a potential obstacle for causal interpretation. Manski argues that a key issue for peer effect studies concerns the direction of information flow between the peers, particularly if multiple individuals receive the information shock. In our setting, this could pose a problem if multiple households in the same neighborhood file warranty claims, even if they do so at different times. As this would make some peers both information senders and information receivers, it would be difficult to determine which warranty claim by which peer had a causal impact on the purchase decisions of their peers.

4.2 Sample Selection Criteria to Address the Reflection Problem

To address the reflection problem, we impose two sample criteria. First, we only include those claims whose neighborhoods have exactly one warranty claim during the whole period. Second, we only include those claims whose neighborhoods do not overlap spatially with the neighborhoods of

[17]Butts (2023) contends that this approach (of examining the distance at which the estimated treatment effect approaches zero, and then remains at zero) is "akin to the pre-trends test in event study regressions," in that it provides a data-driven method for determining whether the identification assumptions are appropriate.

any other warranty claim. By using these sample criteria, we ensure that both the treatment and control group neighbors of a warranty claim are impacted only by that warranty claim.

These sample selection criteria are quite restrictive in that, from an initial population of over 95,000 claims, we exclude nearly 98% of claims. However, without these sample criteria, the causal interpretation of our analysis (and any others like it) is suspect. Therefore, a valuable advantage of our data is that our original population is sufficiently large that, even after we impose these strict criteria, we still have 1,776 warranty claims and their neighborhoods on which to perform our analysis.

4.3 Sample Criteria and Staggered Events in DID

Our sample criteria also address another econometric issue: bias in two-way fixed effects estimation due to staggered events in DID models. A large recent literature has shown that DID models with staggered timing of events when there are heterogeneous effect sizes across treatments can have biased treatment effect estimates.[18] Various solutions have been proposed for this problem, including only using control group units that are "later-treated" or "never-treated".

In our setting, if we were to use our full sample (rather than the carefully chosen sample), then we would face these problems related to staggered event timing. This is because the full sample includes many neighborhoods with multiple warranty claim events, made by multiple neighbors, on multiple dates.

However, as our analysis sample only includes neighborhoods with exactly one warranty claim event affecting households over the entire period, within each warranty claim's neighborhood, there is only a single event over the whole time period. Effectively, for each warranty claim, we run a "sub-experiment" wherein we compare the effect of that unitary warranty claim on treatment group units (inner ring postal codes) against its effect on control group units (outer ring postal codes). There are no other warranty claim events that may affect the treatment status of these units at other dates. In other words, our control group households are effectively "never-treated" units and our treatment group households are only treated once. For these reasons, because of our carefully chosen analysis sample, we are not exposed to the problem of staggered DID analysis.

[18] See Goodman-Bacon (2021); Baker et al. (2022); Cengiz et al. (2019) and many others.

4.4 Sample Criteria and Akerlof (1991)

While the sample selection procedure suggested here, of using *exactly one* claimant, is specifically designed to address various econometric issues (i.e. the reflection problem and staggered DID), it is also advantageous in our setting for substantive economic reasons. In particular, it allows us to more directly test the "cognitive psychology" proposition of Akerlof (1991). This proposition states that a *single* salient peer will have a large effect on the consumption choices of the information receiving peer. Thus, our sample selection criteria of restricting our sample to neighborhoods with *exactly one* information sending warranty claimant allows us to directly test the Akerlof (1991) proposition.

4.5 Using a Less Strict Sample Selection Criteria

In our discussion here, we state that using this strict sample selection criteria provides important econometric advantages (reflection problem and staggered DID) and allows us to directly test the proposition in Akerlof (1991). However, as we note above, this strict sample selection criteria drops approximately 98% of our warranty claim observations. One possible concern that arises as a consequence is that, by focusing on this strict sample selection criteria, we reduce the external validity of our results.

To address this concern, we replicate our tests but with significantly weaker sample selection criteria. As we describe above, our main tests make the strong assumption that there are no other warranty claim events in that neighborhood throughout the whole of our sample (where n = 1,776 claim events, or 2% of our total claim events). As a robustness check, we also run our baseline specification with a weaker sample selection criteria, which includes neighborhoods if there is no other warranty claim in the neighborhood within 8 quarters of the warranty claim event. This results in a regression with 25,483 claim events, or 27% of all claim events, which is significantly larger. As we describe in Section 6.2, our main finding is that the main DID coefficient is still highly significant in this weaker sample selection test, although the magnitude of the DID coefficient is smaller.[19]

[19]This smaller estimated magnitude serves as evidence supporting our argument for selecting a more restricted sample. In this setting, which includes more claim events, multiple claims can occur within the defined circle. As a result, the effects of previous claims may contaminate the estimation of subsequent events and result in a smaller estimate of the DID coefficient.

5 Empirical Methodology

Our primary specification may be represented as follows:

$$Y_{i,c,t} = \beta_1 Inner_{i,c} + \beta_2 Post_{t,c} + \delta Inner_{i,c} \times Post_{t,c} + \gamma_i + \gamma_t + \gamma_c + \gamma_\tau + \varepsilon_{i,c,t}, \quad (1)$$

where $Y_{i,c,t}$ is a measure of the purchases of the claimed brand relative to all other brands in postal code i in the neighborhood of claimant c in the two years before and after the claim filing ($t \in \{0, 1\}$), $Inner_{i,c}$ indicates whether postal code i is within the inner ring for warranty claimant c's postal code, $Post_{t,c}$ indicates whether the two-year time period t occurs before or after claimant c files their warranty claim, γ_i, γ_t, γ_c, and γ_τ are postal code, claim-relative time period, warranty claimant, and warranty claim quarter fixed effects, respectively,[20] and δ is the coefficient of interest.[21] If an idiosyncratic and unrepresentative signal from a peer affects consumption, we predict that a warranty claim event will reduce the subsequent purchases of the claimed brand relative to all other brands for neighbors living in postal codes that are very close to the claimant's postal code relative to neighbors living in slightly further away postal codes, i.e., we predict that δ is negative.

5.1 Dependent Variable

Our main hypothesis predicts that the warranty claim on a specific brand by a neighbor will reduce the sales of the claimed brand relative to the sales of other, unclaimed brands. In our baseline specification in Equation 1 above, we stipulate that our dependent variable should measure the sales of the claimed brand (in the specified time and place) relative to the sales of *all other* unclaimed brands (in that same time and place). There are two possible versions of such an outcome variable: (1) net sales, which is sales of the claimed brand *minus* sales of all unclaimed brands, and (2) market share, which is sales of the claimed brand *divided by* sales of all brands.

We choose to use the net sales measure as the market share ratio measure does not work well for us because, in our setting, it is often missing because of a divide-by-zero problem. Our analysis is run at the postal code-claim-relative time period-level and a large number of postal codes in the neighborhoods in our study (approximately half) have zero transactions of *any* brand in a specific

[20] Another way of describing our analysis, due to the warranty claimant fixed effect we employ, is to note that each neighborhood can be considered as a separate sub-experiment with a single event. Our specification essentially stacks all of these different sub-experiments in a way that is similar to Cengiz et al. (2019); Deshpande and Li (2019).

[21] Given that we implement postal code and claim-relative time period fixed effects, our regression analysis does not estimate β_1 or β_2. We include these two terms in our specification for completeness.

quarter.[22] Because the total sales in these postal code-periods are zero, these observations have undefined market share values and, as such, are dropped from our analysis if we use market share as our dependent variable. For this reason we use the net sales measure (based on subtraction rather than division) as our dependent variable.

5.2 DID Assumptions: Parallel Pre-Trends

In this section, we confirm parallel trends in our outcome variable prior to the warranty claim filing using a dynamic DID specification, with the $t-1$ quarter omitted. We do this to assess the validity of our DID as, if there were non-parallel trends prior to the claim, it would be quite likely that our DID coefficient estimate had a non-parallel trends bias (Cunningham, 2021). We examine four quarters before and after the warranty claim event quarter. Results for these dynamic DID tests are reported in Table 2. This table reports 12 different specifications with varying ring sizes (the width of the treatment and control rings) and separations (the gap between the treatment and control rings).

Table 2 shows that, across *all* of the specifications, we find no significant coefficients in pre-claim quarters, strong evidence consistent with the parallel pre-trends assumption. In addition, Table 2 shows significant negative coefficients, as predicted by our main hypothesis, in many post-claim quarters, especially in quarters 3 and 4 after a claim. We plot the dynamic DID coefficients for the first column of Table 2 in Figure 3. The figure confirms that the claim event has a negative effect on peers' purchasing behavior for the claimed brand right away, though it becomes statistically significant only in quarters 3 and 4. Equally importantly, the figure confirms that there is no difference in pre-claim purchasing behavior between inner and outer ring neighbors, confirming pre-filing parallel trends for our DID analysis.

5.3 Determining Size of Treatment and Control Rings

As described in Section 3.4, we use the data-driven methodology developed in Butts (2023) to define the width of treatment and control rings in our DiD. Briefly, following Butts (2023), we estimate nine treatment effects of the warranty claim events for households in increasingly distant treatment rings (the red rings in Figure 2) relative to households in a tenth control ring beyond all nine

[22]For instance, in our analysis sample, we observe transactions for 6,948 postal code-quarters when we use the net sales (subtraction) measure but, because of the divide-by-zero problem, we only have data for 3,687 postal code-quarters using the market share ratio measure.

treatment rings (the blue ring in Figure 2). We define ring widths for this non-parametric exercise in two ways: setting the width of the ten rings equal and setting the area of the ten rings equal. We report the results of these estimations in Figure 4 and Table 3. The first panels of Figure 4 and Table 3 present our findings using equal width rings and the second panels present the equal area ring findings. Our two main conclusions from these tests are: (1) the DID coefficient in the innermost ring is highly significant and negative, indicating that there is a very large treatment effect for neighbors within the innermost ring, and (2) that the DID coefficient effectively converge to zero at ring 2.

These findings have a number of important implications. First, they are consistent with the key assumption of the hyper-local DID identification strategy, which is that the closer neighbors are to each other, the more likely they are to interact with each other (i.e. to have discussed the warranty claim). Our findings of a highly significant DID coefficient only for ring 1 implies that only the most inner ring neighbors in ring 1 have statistically significant interactions with the warranty claimant.

Second, these results showing that the DID coefficients for rings 2 to ring 9 are all effectively zero indicate that the effect of the warranty claim is essentially constant across all rings from the second innermost ring (ring 2) to the outermost control ring (ring 10). This is because any effect from the treatment in rings 2 to 9 is "differenced out" when it is compared to the control ring in ring 10, resulting in a series of DID coefficients being equal to zero. This finding supports the spatial "local parallel trends" assumption in Butts (2023) that any other possible treatment effect is constant across space, and can thus be "differenced out" in a spatial DID specification.

Third, the estimated DID coefficient is much larger in absolute magnitude for ring 1 in the equal distance specification (where each ring, including ring 1 is 100 m) than for ring 1 in the equal area specification (where ring 1 is 316 m). This is strongly consistent with the argument that the treatment effect (as measured by the magnitude of the DID coefficient) dissipates rapidly as the distance from the warranty claim event increases. For this reason, in the rest of the paper we focus on the equal distance specification, where each ring is 100 m.

5.4 Exclusion Buffers Between Treatment and Control Rings

As described above, the key problem in the rings-based hyper-local identification strategy is that it is often not clear where to draw the rings containing treatment and control units. We cannot observe which neighbors are truly treated (i.e., interact with the sending peer) and untreated (i.e., do not interact with the sending peer). Above, we describe one possible solution to this issue, which

is the Butts (2023) methodology of comparing all red rings to a "last" blue ring.

However, another possible (and somewhat related) solution to this problem is the imposition of exclusion buffers (i.e., gaps in space), between the treatment and control rings as done in Schiff et al. (2023). By imposing an exclusion buffer between the treatment and control rings, the empiricist can increase the likelihood that the treatment and control neighbors are accurately differentiated. The choice of a "buffer zone," however, involves a trade-off. To improve the causal interpretation of any empirical analysis, we must ensure that the treatment and control groups are as similar to each other as possible, across both observable and unobservable characteristics. As such, a very small exclusion buffer (e.g., of 0 m between treatment and control rings) improves the comparability of treatment and control groups (across both observables and unobservables) by reducing the distance between any two households in the treatment and control rings. However, having no "buffer zone" separation between treatment and control rings also increases the likelihood of contaminating our control group by including households who did interact with the claimant and contaminating our treated group by including households who did not interact with the claimant. This lack of geographic separation between the rings may reduce the contrast between the rings and weaken our ability to identify true peer effects.

In this section, we examine this trade-off in our setting by conducting a series of DID tests with differing exclusion buffer sizes, ranging from an exclusion buffer of 0 m to an exclusion buffer of 800 m. All of these tests use a 100 m-wide innermost ring as the treated group (ring 1 in Figure 2) and nine 100 m-wide control rings with steadily increasing exclusion buffers between them.[23]

We show the results of this test in Figure 5 and Table 4. We find that the DID coefficient remains constant across all of the exclusion buffers ranging from 0 m to 800 m. Given the aforementioned advantages of having narrow, directly abutting treatment and control rings and the findings reported here, we choose to use the innermost 100 m-wide ring (ring 1 in Figure 2) as our treatment group and the next 100 m-wide ring (ring 2 in Figure 2) as our control group, with no exclusion buffer, as our specification for the remainder of this paper.

[23]Note that these tests are quite similar in spirit to the Butts (2023) tests we describe above. In those tests the nine treatment rings became increasingly distant from the warranty claim while the control ring remained fixed beyond the ninth treatment ring. In these tests, the treatment group stays fixed as an innermost 100 m-wide ring and, as the exclusion buffer expands, the control rings move increasingly further from the warranty claim.

5.5 Alternative Event Windows

A large literature emphasizes that the timing of consumption for consumer durable goods is "lumpy" (e.g., Attanasio, 2000; Eberly, 1994; Caballero, 1990, 1993; Caplin and Leahy, 2006). As described in this literature, consumption of these kinds of products happens rarely over the lifetime of a consumer, so there may be extended spells during which a consumer (or even a whole neighborhood) does not purchase a consumer durable good. This argument could imply that word-of-mouth peer effects are less effective for consumer durable goods.[24]

Again, there may be a trade-off between the infrequency of consumer durable purchases and salience of recent word-of-mouth information for our tests when selecting the amount of time before and after the claim events that we study. To choose the most appropriate time window, we run our baseline analysis using eight different time windows around claims. Specifically, we perform our classic DID regression, where the time periods before and after the claim event range from 1 quarter to 8 quarters. The results of these tests are displayed in Table 5 and Figure 6. Both the table and the figure show that the main DID coefficient is significant and negative (as predicted) across all of the eight time windows. In addition, the magnitude of the DID coefficient increases monotonically (in absolute value) as the length of the event window increases from 1 quarter to 8 quarters.

These results imply that the magnitude of the peer effect spillover in consumption gets larger as the time window around the claim event expands. This is consistent with the notion that durable good consumption is "lumpy." The receiving peer, even if they receive the information shock when the warranty claimant files, may only use the information when they choose to make a durable good purchase. As we expand our time window around the claim event, the proportion of peers making purchase decisions for consumer durable goods increases. Though the recency of the claim event does not seem to weaken the word-of-mouth peer effect, in order to balance the effects of the infrequency of durable good purchases and the potential salience of recent information transfers, we choose the four quarter (one year) time window for our baseline analyses.

[24] Dranove and Jin (2010), in its literature survey of various kinds of quality disclosure mechanism (p. 937), states that "word-of-mouth (effects) are ... of limited value when products are infrequently purchased".

6 Results

6.1 Baseline Findings

In the previous sections, we carefully explored the parameters of the hyper-local DID framework we use for our peer effects analysis. First, we defined our sample such that each inner and outer ring neighbor only experiences one peer warranty claim event. Next, we chose the width of rings and separation between inner and outer rings such that inner ring neighbors are within 100 m of the claimant and outer ring neighbors are between 100 and 200 m away from the claimant. Finally, we chose to use the four quarters (one year) before and after the warranty claim as the time period during which we conduct our analyses.

Using the above parameters, we find that, in the four quarters after a warranty claim, relative to the four quarters before the claim, the inner ring neighbors of the claimant make 1.6 fewer net purchases of the claimed brand, when compared to outer ring neighbors. Given that households in the typical postal code in our analysis sample purchase 21.8 consumer durable goods in one year, 1.6 fewer net purchases of the warranty claim manufacturer's goods amounts to about 7% of consumer durable goods being purchased from a different manufacturer due to the warranty claim, an economically meaningful shift in the consumption behavior of a claimant's peers. This reduction in the net purchases of the claimed brand (in comparison to other brands) is highly statistically significant ($p < 0.01$). Given our thorough exploration of the parameters of the empirical exercise, we know that this reduction is robust to various perturbations. Overall, our analyses confirm that a warranty claimant's idiosyncratic and unrepresentative product quality information shared via word-of-mouth networks has significant effects on the consumption behavior her neighbors.

6.2 Easing Constraints on Analyzed Claims

To make causal inference possible, in our baseline tests above, we define our analysis sample such that there are no other warranty claims in the neighborhood. As discussed in Section 4, there are two separate econometric reasons for imposing this constraint on our sample, addressing the reflection problem as well as the issue of staggered events for DID estimation.

Even though our baseline specification above is our preferred specification, the restrictive selection criteria result in only a small fraction (approx 2%) of all warranty claim events being used in our analyses, raising the concern that our results may be unique to this subsample. In this section, as a robustness check, we conduct our baseline specification on a more loosely-constrained sample. For

this larger sample, we include warranty claims if there are no previous claims in the neighborhood within 4 quarters of the warranty claim event.

We report the results of these tests in Table 6. The main finding of this table is that the peer effects DID coefficient is still highly significant and negative (as was the case with the baseline specification). As expected, the sample size of this new test is significantly larger than the more constrained sample in our baseline specification, because many more neighborhoods now fit into the sample selection criteria. This larger sample includes 27% of all warranty claim events. The magnitude of the coefficient is smaller than in the baseline specification, which is related to the fact that a prior warranty claim (5 or more quarters prior to the event date) may be contaminating our estimated effect. The receiving peer could be reacting to either, or both, of the claim in our test, or any unobserved claims 5 or more quarters prior to the event date. Such previous claims are likely to attenuate the response we measure in our specification.

6.3 Long Distance Relationships that are Not Based on Proximity

In our results from the Butts (2023) procedure above, we document that the treatment effect is very large and significant in the innermost ring but then dissipates to zero in the remaining concentric rings. These findings strongly suggest that proximity is a crucially important component of the estimated treatment effect, i.e., that very proximate geographic neighbors (within 100 m) interact with sending peer (warranty claimant), while slightly more distant neighbors (beyond 100 m) do not interact with them.

However, it is also possible that neighbors interact with each other using mechanisms that do not vary by distance at the scales examined in our study. For example, it is possible that various neighbors in our study have "long distance" relationships with the information sender, such as relationships over internet networks (e.g., via Facebook). This example of "long distance relationship" treatments is taken from Bailey et al. (2022), who examine consumption peer effects among Facebook friends. That paper's key finding is that the estimated treatment effect is still significant for more distant Facebook friends who live farther away. Specifically, these more distant friends in that paper live in different U.S. zipcodes, which are orders of magnitude larger in population than the Canadian six-digit postal codes we study.

The Butts (2023) approach that we implement above helps address this issue. The DID coefficients we estimate for increasingly more distant treatment rings allow us to assess if there is a treatment effect that varies across the distances we study. And, given the convergence to zero of

the effect from the second-innermost ring outwards, we find that there is, indeed a treatment effect that varies at the distances we study. Moreover, any long-distance relationship between the sending peer and receiving peers that is constant at the distances we study are differenced out using our DID methodology. As such, we are confident that, in our main DID analysis, we are identifying the effects on neighbors' purchases of a warranty claim effect that declines over the distances examined in our study.

7 Idiosyncratic Peer Experience vs. Systematic Product Quality Information

We have already shown that individuals behave as described by the "cognitive psychology" view posed by Akerlof (1991), wherein a single salient individual's warranty claim has an effect on the consumption choices of her peers. There is, however, a potential concern that the information shared by the sending peer (warranty claimant) to the receiving peer contains information not just based on her own idiosyncratic and unrepresentative experience, but also some other, more representative, information about product quality. In other words, the concern is that we are mislabeling information shared by the warranty claimant as entirely idiosyncratic and unrepresentative of product quality, as needed to assess the validity of the Akerlof (1991) proposition. If that were true, then the altered purchasing behaviors of the receiving peers (the neighbors) that we observe in our tests could be driven by the representative (i.e., accurate) product quality information sent by the sending peer.

To address this concern, in this section, we examine the effect of the warranty claim on neighbors' purchasing decisions accounting for three alternative and more representative measures of product quality: a direct measure of product quality (aggregate warranty claim rate); an indirect measure of product quality (brand market share); and, a direct measure of warranty quality (days to repair).[25] If neighbors' purchases are affected by the component of the claimant's information that is correlated with any of these representative measures of quality, then the peer effect of the warranty claim (i.e. the DID coefficient) should become smaller and/or less significant when we include these measures as additional RHS variables in our baseline specification.

[25]These are some typical examples of quality assurance mechanisms for consumer durable goods, as highlighted in Dranove and Jin (2010). That paper indicates that such quality assurance mechanisms are preferred to word-of-mouth quality disclosure (such as our warranty claimant) as these mechanisms "provide more precise and comparable information than word-of-mouth" (p. 939).

As we show below, however, when we include any or all of these three systematic measures of quality (either individually or combined) as additional explanatory variables in our specification, the DID coefficient (measuring the effect of the warranty claim on neighbor's purchasing decisions) remains unchanged. This result indicates that the idiosyncratic and unrepresentative component of the sending peer's (warranty claimant) signal is driving the receiving peers' purchasing decisions.

7.1 Aggregate Claim Rate

Our direct measure of product quality is the aggregate warranty claim rate for each brand. We define it as the total number of claims made against that brand divided by the total number of sales of that brand's products for the claimed product category in the two years before each focal claim event. Because of the very large size of our data, we argue that the aggregate warranty claim rate measure provides an accurate and direct measure of the true underlying product quality, when product quality is defined based on the fraction of products that break down sufficiently to require warranty claims.

To exclude any information about aggregate claim rate from the information shared by the warranty claimant, we run our baseline specification (Equation 1) while controlling for the claim rate for the brand in the two years before the warranty claim, described above.[26] In order to run the analysis with the claim rate control variable, we have to take away two sets of fixed effects, the warranty claim event fixed effect and the postal code fixed effect, and replace them with an inner ring fixed effect.[27]

We provide the results for this analysis in column (2) of Table 7. We find that the impact of the warranty claim on the claimant's neighbors' purchases (i.e. the DID coefficient) does not change when we control for the claimed brand's aggregate claim rate. This can be seen when comparing the DID coefficient in column (2) which includes the control, to the DID coefficient estimated in column (1) without controls. This finding indicates that the effect of the warranty claim on neighbors' purchasing decisions is not driven by any component of the claimant's signal correlated with this direct measure of product quality, aggregate warranty claim rate.

[26] As a robustness check, we run these tests using overall brand-level claims-to-sales ratios and the results are qualitatively similar.

[27] This easing of fixed effects necessarily means that the analyses in this section should not be used to establish causal identification.

7.2 Brand Market Share

Next, we test our main effect while controlling for an indirect measure of product quality, brand market share.[28] There is a substantial literature in marketing (e.g., Hellofs and Jacobson, 1999) as well as international trade (e.g., Khandelwal, 2010) which asserts that, in settings where quality is unobservable by consumers, consumers attribute higher levels of quality to brands with higher market shares, conditional on price.[29]

We measure the market share of the claimed product in the same way as for aggregate claim rate. In particular, we define it as the number of purchases of the brand's products divided by the total number of purchases for the claimed product category in the two years before each focal claim event. We use this measure as a control when we run our baseline specification, again in the same manner as described in the previous subsection.

The results for this analysis are available in column (3) of Table 7. The impact of the warranty claim on the claimant's neighbors' purchases (the DID coefficient) does not change when we control for the claimed brand's product market share, as can be seen when comparing the DID coefficient in column (3) to the DID coefficient estimated in column (1) without controls.[30] We interpret this as indicating that the effect of the warranty claim on neighbors' purchasing decisions is not driven by the claimant's signal having any information correlated with our indirect measure of product quality, product market share.

7.3 Warranty Claim Process Length

Finally, we test the effect of an unrepresentative peer's information while controlling for a direct measure of warranty quality, days required to repair the good. The days to repair information is a part of the claimant's experience that is most likely to be representative of some true measure of product quality, so it is important that we test whether our main effect is driven by it.

To run this analysis, we measure days to repair as the number of days taken to repair the good,

[28] Given that the retail chain providing us with data is one of the top four retailers of consumer durable goods in Canada, we argue that we can representatively measure the market share for all brands sold in the total Canadian market.

[29] The argument posed in those literatures is that the consumer believes that if a large number of other consumers are already purchasing that brand, then the unobserved quality of that brand must also be high.

[30] As with the claims rate analyses, we run these tests using overall brand-level product market share and find qualitatively similar results.

starting from the date of the initial warranty claim to the date of the final resolution of the claim. To test whether the peer effect of the warranty claim on neighbors' purchasing behavior is driven by potentially representative information about days required to repair, we include this measure as a control on our baseline regression, as described above.

Our findings are reported in column (4) of Table 7. The DID coefficient for the peer effect when controlling for days to repair is nearly identical to the DID coefficient in column (1), which has no controls. This indicates that the effect of the warranty claim on neighbors' purchasing decisions is not driven by information shared by the claimant about warranty quality, as measured using days to repair.

So far in this section, we have shown that the peer effect we document is not correlated with any of the three product quality measures we explored: brand claim rate, brand market share, and days required to repair. In column (5) of Table 7, we use all three of these product quality measures as controls for our baseline regression, to test whether our peer effect is perhaps driven by some combination of these three objective measures of product quality transmitted by the warranty claimant to her neighbors. As with our tests adding the controls separately, this test shows that our peer effect (as measured by the DID coefficient) remains unchanged by the addition of all three controls simultaneously, as evidenced by the DID coefficient reported in column (5) being effectively the same as the DID coefficient reported in column (1) of the table, wherein we had no additional controls.

Put together, these findings allay concerns that the peer effect of the warranty claim on the receiving peers' purchasing behavior may be driven by some component of the sending peer's signal that is representative of true product quality. The DID coefficient (reflecting the impact of the warranty claim on the receiving peers) remains large, highly significant, and effectively unchanged even when we account for these three product quality measures. This indicates that the idiosyncratic and unrepresentative component of the sending peer's signal drives the receiving peer's consumption choices.

8 Peer Effect Strength across Neighborhood Characteristics

Neighborhoods differ in various ways which may impact the strength of linkages between neighbors (Chetty et al., 2022) and, as a result, peer effects could vary across different neighborhoods based on these differences. In this section, we examine the extent to which three measurable neighborhood

characteristics — average income, population density, and immigrant proportion — determine the peer effect of a warranty claim on the claimant's neighbors' purchasing behavior. Overall, we find that these neighborhood characteristics play a significant role in the observed peer effect.

8.1 Measuring Variation Across Neighborhood Characteristics

Our neighborhood characteristics are at the dissemination area (DA) level. The DA level is the smallest level of geography at which census data is made available by Statistics Canada. While not as small as the postal codes that we use as our unit of analysis, DAs are still quite small, with the median DA including 550 people. We use these DA-level census data for neighborhood-level demographic information.

To study how neighborhood characteristics affect the peer effects of warranty claims, we compare the estimated coefficients for our baseline specification for quintiles of our sample based on the distribution of each neighborhood characteristic. Effectively, we run our hyper-local spatial difference-in-differences analysis for five subsamples of warranty claims for each characteristic, based on the quintile of the characteristic in which the claimant's DA belongs. We provide summary statistics of these quintiles in Table 1e.

8.2 Neighborhood Average Income

Recent work from Chetty et al. (2022), using Facebook data, examines "friending rates" across SES (Social and Economic Status) percentiles. This "friending" study provides evidence that as SES declines, the number of friendships developing from neighborhood interactions increases. In other words, this study shows that neighborhoods with lower SES (e.g., neighborhoods with lower income households) are more likely to have friends based on peer interactions from within the neighborhood. Based on Chetty et al. (2022), we hypothesize that our peer effects should be larger in neighborhoods with lower average income, as word-of-mouth information spreads more effectively within neighborhoods with stronger social networks.

We use Canadian census data on average income in a neighborhood to test the validity of our hypothesis. Our results are reported in Table 8 and Figure 7. As is graphically evident in Figure 7, we find a significantly larger and more negative consumption peer effect for the lowest income quintile (quintile 1) than for all other quintiles. And, for all but the highest quintile, there is a clear monotonic decrease in the magnitude of the peer effect as average income increases. These results are consistent with the hypothesis, based on Chetty et al. (2022), that the magnitude of

neighborhood peer effects should be the largest for individuals in the lowest SES neighborhoods.

8.3 Neighborhood Population Density

An extensive literature across many disciplines (e.g., urban economics, urban sociology, and geography) has examined how variation in population density impacts the strength of the relationship between neighbors. In general, two contradictory hypotheses have been proposed. On the one hand, it is argued that increased population density should generate more social interactions, and thus larger peer effects, because increased density increases the opportunity to interact with more peers (Boessen et al., 2018). On the other hand, it is argued that increased density causes increased interpersonal frictions, offers other alternatives for establishing social networks, and generates other various negative externalities (e.g., congestion externalities), which, in aggregate, lower the quality of neighborhood-based social interactions, and thus lower the impact of such peer effects (Brueckner and Largey, 2008; French et al., 2014; Skjaeveland and Garling, 1997; Mouratidis and Poortinga, 2020). As such, in our setting, it is an empirical question whether population density should increase the magnitude of the effect of warranty claims on neighbors' purchasing behavior or decrease it.

We examine how population density affects our peer effect using DA-level population density data. We are unable to run our analysis on the lowest density quintile as our sample selection criteria mechanically excludes that quintile.[31] Our results for this specification are reported in Table 9 and Figure 8. As can be seen in Figure 8, the peer effect estimate for the least dense quintile (i.e., quintile 2) is significantly more negative than all other quintiles. Moreover, the peer effect is monotonically decreasing with population density. This finding is consistent with the hypothesis that social interactions are stronger in less dense neighborhoods, because peer interactions in more dense neighborhoods are impeded by issues such as interpersonal frictions, negative congestion externalities, and availability of alternative social networks.

8.4 Neighborhood Immigrant Proportion

Finally, we examine whether the fraction of immigrants within a neighborhood affects the impact of word-of-mouth peer effects such as ours. A number of factors could drive heterogeneity in these neighborhood peer effects based on the immigrant share of a neighborhood. For instance, Ray and Preston (2009) posits that immigrants may be more socially isolated from their non-immigrant

[31] Recall that we choose neighbors in postal codes within 200 m of the claimant's postal code for our inner and outer rings. Rural postal codes are much larger than that and, as a result, are mechanically excluded from our analyses.

neighbors in the neighborhood. Such socially-driven divisions within a neighborhood would make peer effects like the one we study here less powerful.

Again, we divide the neighborhoods in our study into five quintiles based on the fraction of immigrants in the population, as extracted from our DA-level Census data. We run our quintile-based tests for immigrant fraction and report our findings in Table 10 and Figure 9. In the figure, we observe a general trend that the magnitude of the peer effect declines as the fraction of immigrants increases. While the trend is not as dramatic as the trends we observe for neighborhood average income or population density, it seems to be economically relevant, with the neighborhoods with the highest immigrant proportions reporting a 40% weaker peer effect than the lowest immigrant quintile neighborhoods. This weaker peer effect is consistent with neighborhood social networks being weaker where there are greater proportions of immigrants, who may be more socially isolated within their geographic neighborhood.

9 Conclusion

Akerlof (1991) in his Ely Lecture, describes a vignette where a single salient, but possibly unrepresentative, peer who has had a negative quality experience of a brand can deter another peer from purchasing that brand. Akerlof (1991) labels this vignette as reflecting the "cognitive psychology" view, which, he argues is very different from the predictions of standard economics. The standard economics view suggests that individuals should not be influenced by a single peer, with unrepresentative information.

In this paper, we provide causal evidence testing the validity of this vignette. Our setting examines a single neighbor in a neighborhood making a warranty claim on a durable of a specific brand. Warranty claims can be considered as unrepresentative information regarding the quality of the product, because only 1.5% of all durable goods sold result in a warranty claim. Our main finding is that a negative experience with a particular brand significantly reduces the future consumption of that brand's products, by that person's geographic peers. Our paper thus provides important new evidence that is consistent with the "cognitive psychology" view of Akerlof (1991), but that is inconsistent with the standard economics view.

References

Sumit Agarwal, Vyacheslav Mikhed, and Barry Scholnick. Peers' income and financial distress: Evidence from lottery winners and neighboring bankruptcies. *The Review of Financial Studies*, 33(1):433–472, 2020.

Sumit Agarwal, Wenlan Qian, and Xin Zou. Thy Neighbor's Misfortune: Peer Effect on Consumption. *American Economic Journal: Economic Policy*, 13(2):1–25, May 2021. ISSN 1945-7731, 1945-774X. doi: 10.1257/pol.20170634. URL `https://pubs.aeaweb.org/doi/10.1257/pol.20170634`.

George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

George A Akerlof. Procrastination and obedience. *The American Economic Review*, 81(2):1–19, 1991.

Elliot Anenberg and Edward Kung. Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures. *American Economic Review*, 104(8):2527–2551, August 2014.

Desmond Ang. The Effects of Police Violence on Inner-City Students. *The Quarterly Journal of Economics*, 136(1):115–168, December 2020. ISSN 0033-5533, 1531-4650. doi: 10.1093/qje/qjaa027.

Peter Arcidiacono, Paul B Ellickson, Carl F Mela, and John D Singleton. The competitive effects of entry: Evidence from supercenter expansion. *American Economic Journal: Applied Economics*, 12(3):175–206, 2020.

Mohammad Arzaghi and J Vernon Henderson. Networking off madison avenue. *The Review of Economic Studies*, 75(4):1011–1038, 2008.

David Atkin, Amit K Khandelwal, and Adam Osman. Exporting and firm performance: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 132(2):551–615, 2017.

Orazio P Attanasio. Consumer durables and inertial behaviour: Estimation and aggregation of (s, s) rules for automobile purchases. *The Review of Economic Studies*, 67(4):667–696, 2000.

Michael Bailey, Drew Johnston, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. Peer Effects in Product Adoption. *American Economic Journal: Applied Economics*, 14(3):488–526, July 2022. ISSN 1945-7782, 1945-7790. doi: 10.1257/app.20200367.

Andrew C Baker, David F Larcker, and Charles CY Wang. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2):370–395, 2022.

Nathaniel Baum-Snow, Nicolas Gendron-Carrier, and Ronni Pavan. Local productivity spillovers. *American Economic Review*, 114(4):1030–1069, 2024.

Patrick Bayer, Kyle Mangum, and James W. Roberts. Speculative Fever: Investor Contagion in the Housing Bubble. *American Economic Review*, 111(2):609–651, February 2021. ISSN 0002-8282. doi: 10.1257/aer.20171611.

Patrick Bayer, Marcus D Casey, W Ben McCartney, John Orellana-Li, and Calvin S Zhang. Distinguishing causes of neighborhood racial change: A nearest neighbor design. Technical report, National Bureau of Economic Research, 2022.

Marianne Bertrand and Adair Morse. Trickle-down consumption. *Review of Economics and Statistics*, 98(5):863–879, 2016.

Adam Boessen, John R Hipp, Carter T Butts, Nicholas N Nagle, and Emily J Smith. The built environment, spatial scale, and social networks: Do land uses matter for personal network structure? *Environment and Planning B: Urban Analytics and City Science*, 45(3):400–416, 2018.

Tessa Bold, Selene Ghisolfi, Frances Nsonzi, and Jakob Svensson. Market access and quality upgrading: Evidence from four field experiments. *American Economic Review*, 112(8):2518–2552, 2022.

Branko Bošković, Sacha Kapoor, Agnieszka Markiewicz, and Barry Scholnick. Add-on pricing over regional business cycles: Evidence from extended warranties. *International Economic Review*, 2024.

Jeffrey R Brown, Zoran Ivković, Paul A Smith, and Scott Weisbenner. Neighbors matter: Causal community effects and stock market participation. *The Journal of Finance*, 63(3):1509–1531, 2008.

Jan K Brueckner and Ann G Largey. Social interaction and urban sprawl. *Journal of urban economics*, 64(1):18–34, 2008.

Leonardo Bursztyn, Florian Ederer, Bruno Ferman, and Noam Yuchtman. Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, 82(4):1273–1301, 2014.

Matias Busso and Sebastian Galiani. The causal effect of competition on prices and quality: Evidence from a field experiment. *American Economic Journal: Applied Economics*, 11(1):33–56, 2019.

Kyle Butts. Jue insight: Difference-in-differences with geocoded microdata. *Journal of Urban Economics*, 133:103493, 2023.

Ricardo J Caballero. Expenditure on durable goods: a case for slow adjustment. *The Quarterly Journal of Economics*, 105(3):727–743, 1990.

Ricardo J Caballero. Durable goods: An explanation for their slow adjustment. *Journal of Political Economy*, 101(2):351–384, 1993.

John Y. Campbell, Stefano Giglio, and Parag Pathak. Forced Sales and House Prices. *The American Economic Review*, 101(5):2108–2131, 2011.

Andrew Caplin and John Leahy. Equilibrium in a durable goods market with lumpy adjustment. *Journal of Economic Theory*, 128(1):187–213, 2006.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The Effect of Minimum Wages on Low-Wage Jobs*. *The Quarterly Journal of Economics*, 134(3):1405–1454, August 2019. ISSN 0033-5533. doi: 10.1093/qje/qjz014.

Raj Chetty, Matthew O. Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B. Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob, Drew Johnston, Martin Koenen, Eduardo Laguna-Muggenburg, Florian Mudekereza, Tom Rutter, Nicolaj Thor, Wilbur Townsend, Ruby Zhang, Mike Bailey, Pablo Barberá, Monica Bhole, and Nils Wernerfelt. Social capital II: Determinants of economic connectedness. *Nature*, 608(7921):122–134, August 2022. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-022-04997-3.

Scott Cunningham. *Causal inference: The mixtape*. Yale university press, 2021.

Janet Currie and Erdal Tekin. Is there a link between foreclosure and health? *American Economic Journal: Economic Policy*, 7(1):63–94, 2015.

Janet Currie, Stefano DellaVigna, Enrico Moretti, and Vikram Pathania. The Effect of Fast Food Restaurants on Obesity and Weight Gain. *American Economic Journal: Economic Policy*, 2(3): 32–63, August 2010. ISSN 1945-7731. doi: 10.1257/pol.2.3.32.

Janet Currie, Michael Greenstone, and Enrico Moretti. Superfund cleanups and infant health. *American Economic Review*, 101(3):435–441, 2011.

Janet Currie, Lucas Davis, Michael Greenstone, and Reed Walker. Environmental health risks and housing values: evidence from 1,600 toxic plant openings and closings. *American Economic Review*, 105(2):678–709, 2015.

Giacomo De Giorgi, Anders Frederiksen, and Luigi Pistaferri. Consumption network effects. *The Review of Economic Studies*, 87(1):130–163, 2020.

Manasi Deshpande and Yue Li. Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy*, 11(4):213–248, November 2019. ISSN 1945-7731. doi: 10.1257/pol.20180076.

David Dranove and Ginger Zhe Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–963, 2010.

Esther Duflo and Emmanuel Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 118(3):815–842, 2003.

Francesco D'Acunto, Alberto G Rossi, and Michael Weber. Crowdsourcing peer information to change spending behavior. *Journal of Financial Economics*, 157:103858, 2024.

Janice C Eberly. Adjustment of consumers' durables stocks: Evidence from automobile purchases. *Journal of Political Economy*, 102(3):403–436, 1994.

Paul B Ellickson and Paul LE Grieco. Wal-mart and the geography of grocery retailing. *Journal of Urban Economics*, 75:1–14, 2013.

Sarah French, Lisa Wood, Sarah Alexandra Foster, Billie Giles-Corti, Lawrence Frank, and Vincent Learnihan. Sense of community and its association with the neighborhood built environment. *Environment and Behavior*, 46(6):677–697, 2014.

Duncan Sheppard Gilchrist and Emily Glassberg Sands. Something to talk about: Social spillovers in movie consumption. *Journal of Political Economy*, 124(5):1339–1382, 2016.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.

Michael Greenstone, Richard Hornbeck, and Enrico Moretti. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of political economy*, 118(3): 536–598, 2010.

Arpit Gupta. Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults. *The Journal of Finance*, 74(5):2249–2301, 2019. ISSN 1540-6261. doi: 10.1111/jofi.12821.

Arpit Gupta, Stijn Van Nieuwerburgh, and Constantine Kontokosta. Take the q train: Value capture of public infrastructure projects. *Journal of Urban Economics*, 129:103422, 2022.

Isaac Hacamo and Kristoph Kleiner. Competing for talent: Firms, managers, and social networks. *The Review of Financial Studies*, 35(1):207–253, 2022.

Bing Han, David Hirshleifer, and Johan Walden. Visibility bias in the transmission of consumption beliefs and undersaving. *The Journal of Finance*, 78(3):1647–1704, 2023.

Rawley Z Heimer. Peer pressure: Social interaction and the disposition effect. *The Review of Financial Studies*, 29(11):3177–3209, 2016.

Linda L Hellofs and Robert Jacobson. Market share and customers' perceptions of quality: when can firms grow their way to higher versus lower quality? *Journal of marketing*, 63(1):16–25, 1999.

Harrison Hong, Jeffrey D Kubik, and Jeremy C Stein. Social interaction and stock-market participation. *The Journal of Finance*, 59(1):137–163, 2004.

Harrison Hong, Jeffrey D Kubik, and Jeremy C Stein. Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers. *The Journal of Finance*, 60(6):2801–2824, 2005.

Rhiannon Jerch, Panle Jia Barwick, Shanjun Li, and Jing Wu. The impact of road rationing on housing demand and sorting. *Journal of Urban Economics*, 140:103642, 2024.

Amit Khandelwal. The long and short (of) quality ladders. *The Review of Economic Studies*, 77 (4):1450–1476, 2010.

Theresa Kuchler and Johannes Stroebel. Social finance. *Annual Review of Financial Economics*, 13:37–55, 2021.

Peter Kuhn, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn. The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery. *The American Economic Review*, 101(5):2226–2247, 2011. URL `http://www.jstor.org/stable/23045637`.

Leigh Linden and Jonah E Rockoff. Estimates of the Impact of Crime Risk on Property Values from Megan's Laws. *American Economic Review*, 98(3):1103–1127, May 2008. ISSN 0002-8282. doi: 10.1257/aer.98.3.1103.

Charles F. Manski. Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3):531, July 1993. ISSN 00346527. doi: 10.2307/2298123.

Gonzalo Maturana and Jordan Nickerson. Teachers teaching teachers: The role of workplace peer effects in financial decisions. *The Review of Financial Studies*, 32(10):3920–3957, 2019.

W Ben McCartney, John Orellana-Li, and Calvin Zhang. Political polarization affects households' financial decisions: Evidence from home sales. *The Journal of Finance*, 79(2):795–841, 2024.

Enrico Moretti. Social Learning and Peer Effects in Consumption: Evidence from Movie Sales. *The Review of Economic Studies*, 78(1):356–393, 2011.

Kostas Mouratidis and Wouter Poortinga. Built environment, urban vitality and social cohesion: Do vibrant neighborhoods foster strong communities? *Landscape and Urban Planning*, 204:103951, 2020.

Paige Ouimet and Geoffrey Tate. Learning from coworkers: Peer effects on individual investment decisions. *The Journal of Finance*, 75(1):133–172, 2020.

Veronika K. Pool, Noah Stoffman, and Scott E. Yonker. The people in your neighborhood: Social interactions and mutual fund portfolios. *The Journal of Finance*, 70(6):2679–2732, December 2015. ISSN 00221082. doi: 10.1111/jofi.12208.

Brian Ray and Valerie Preston. Are immigrants socially isolated? an assessment of neighbors and neighboring in canadian cities. *Journal of International Migration and Integration/Revue de l'integration et de la migration internationale*, 10:217–244, 2009.

Stuart S Rosenthal and William C Strange. Geography, industrial organization, and agglomeration. *review of Economics and Statistics*, 85(2):377–393, 2003.

Nathan Schiff, Jacob Cosman, and Tianran Dai. Delivery in the city: Differentiated products competition among new york restaurants. *Journal of Urban Economics*, 134:103509, 2023.

Oddvar Skjaeveland and Tommy Garling. Effects of interactional space on neighbouring. *Journal of environmental psychology*, 17(3):181–198, 1997.

Figures and Tables

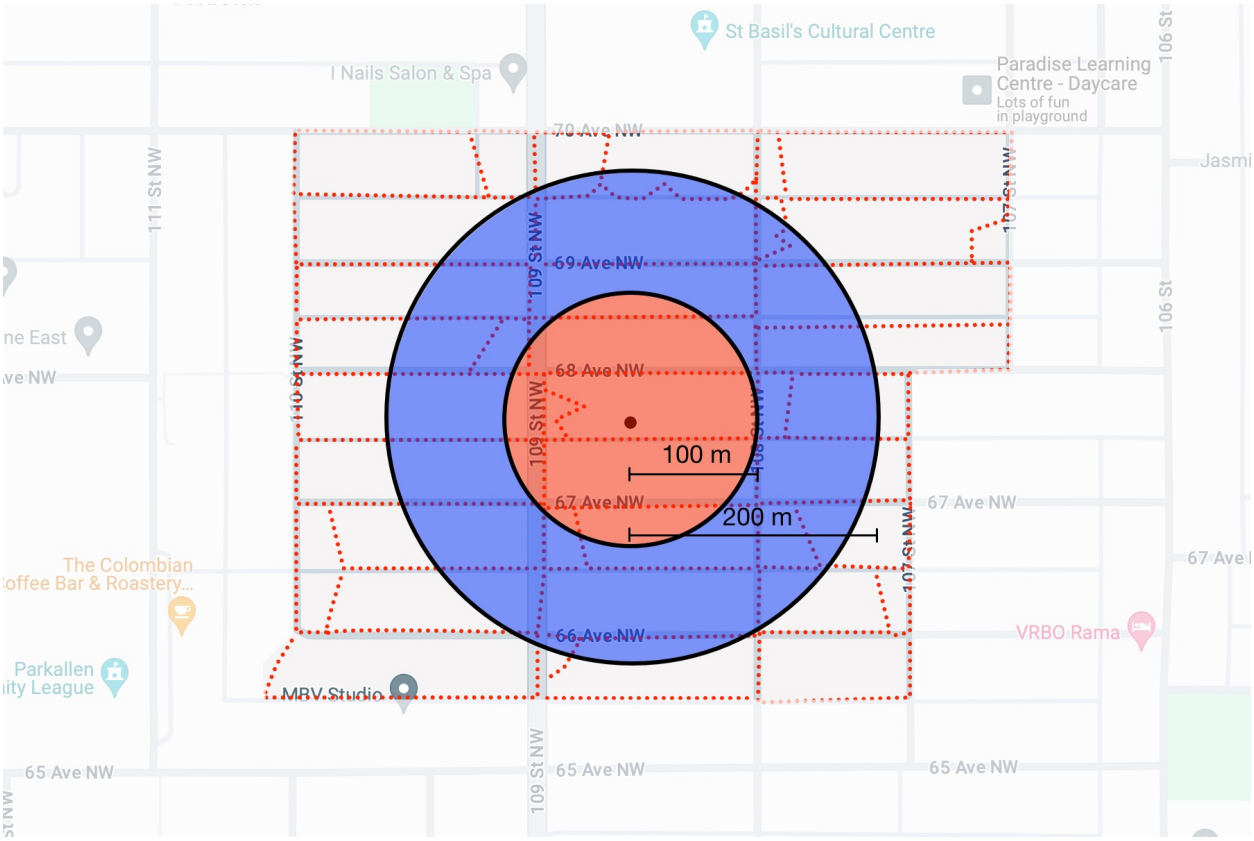


Figure 1: Illustration of Hyper-local Difference-in-Differences Methodology

This figure illustrates the assignment of postal codes to treatment and control groups for the hyper-local difference-in-differences methodology used in this paper. For each warranty claim, an inner circle (red circle with a 100 m radius in this illustrative figure) is drawn around the centroid of the postal code of the claim and all households in postal codes with centroids within the area of the inner (red) circle are in the treatment group (excluding the claimant household). Similarly, an outer ring (blue circle with a 100 m width between 100 m and 200 m from the claimant postcode centroid in this illustrative example) is drawn around the centroid of the postal code of the claim and all households in postal codes with centroids within the area of the outer (blue) ring are in the control group.

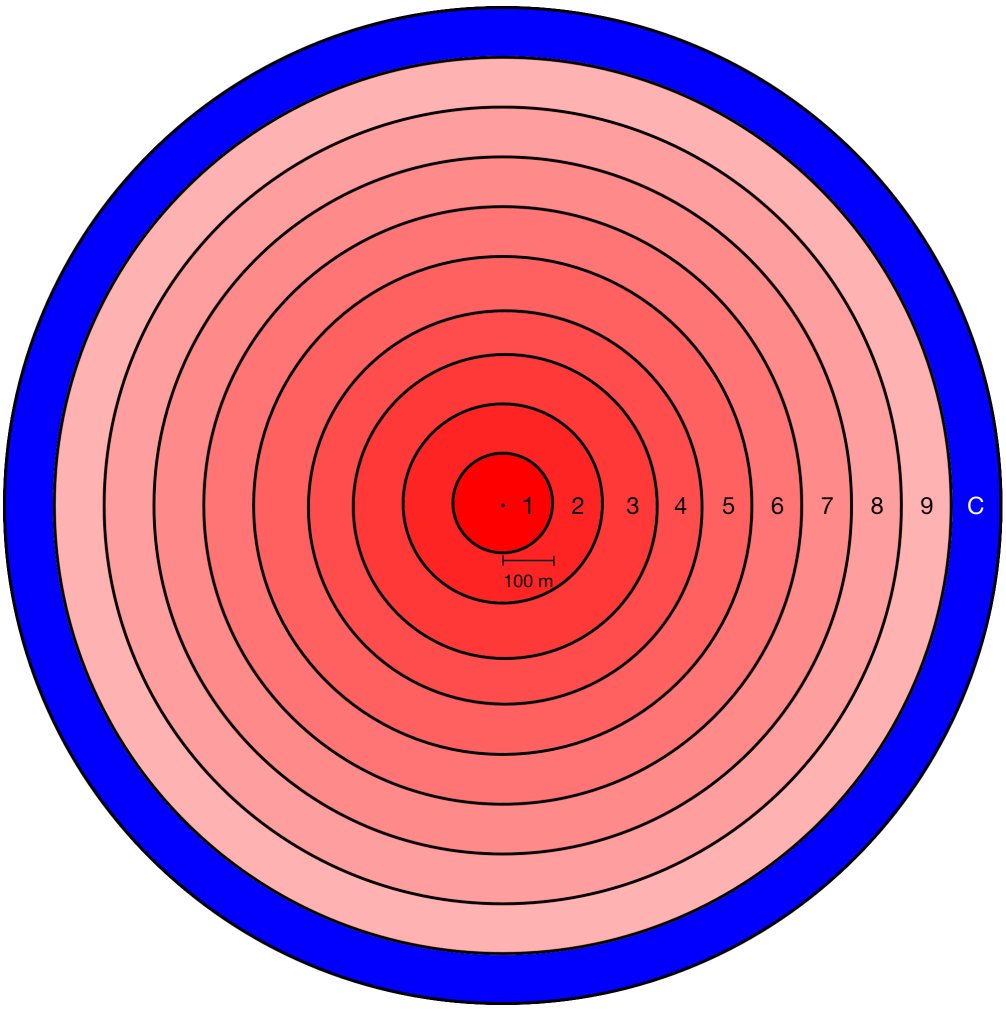


Figure 2: Illustration of Ring-Based Non-Parametric Butts (2023) Methodology
This figure illustrates the estimation methodology from Butts (2023) used in this paper to non-parametrically assess the effect of a warranty claim on neighbors. For postal codes in each of the nine 100 m-wide red concentric rings around the center point, we estimate the change in the outcome variable from before to after the warranty claim and compare the change to the change in the outcome variable for postal codes in the 100 m-wide outermost blue ring.

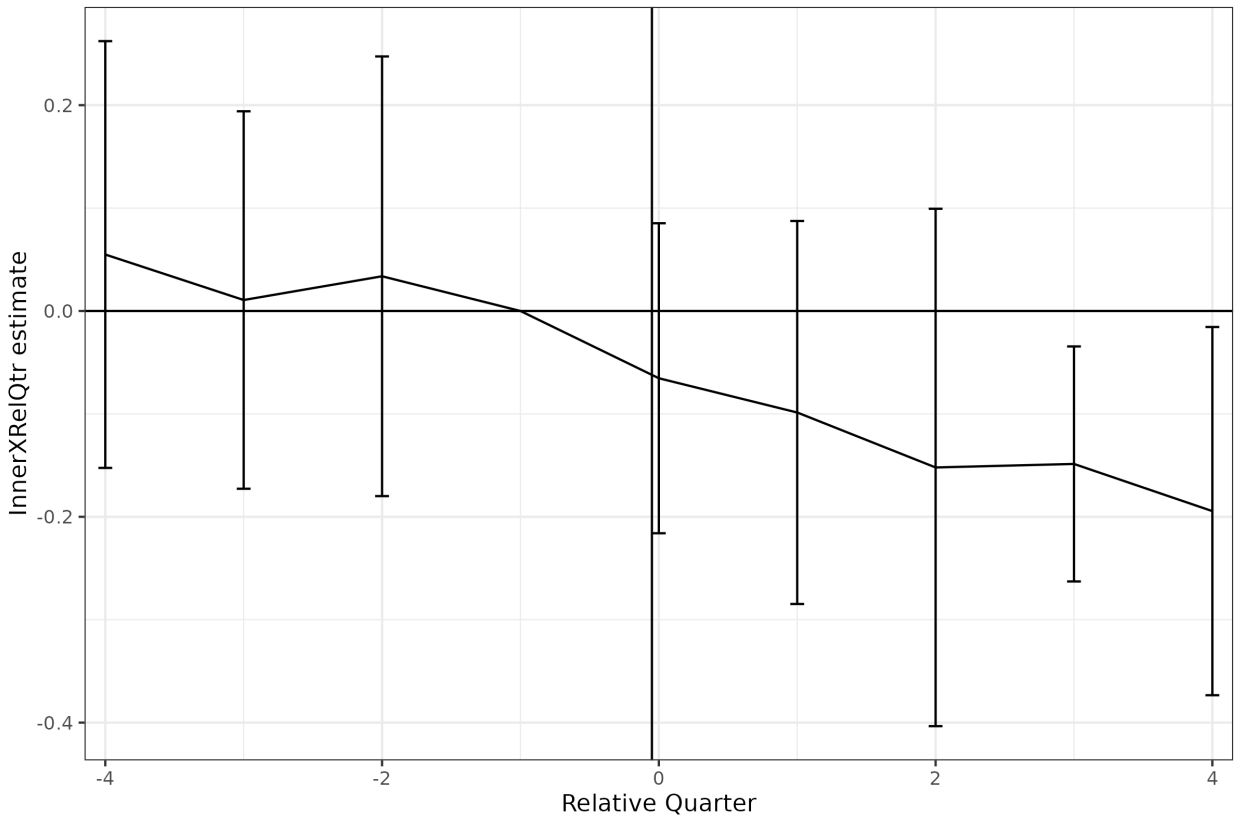
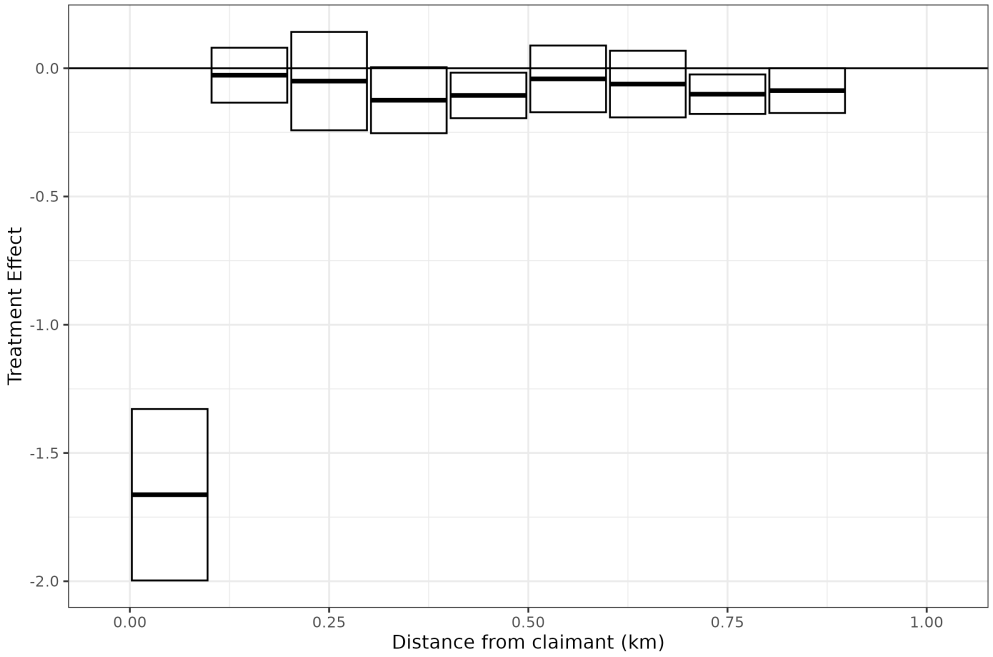
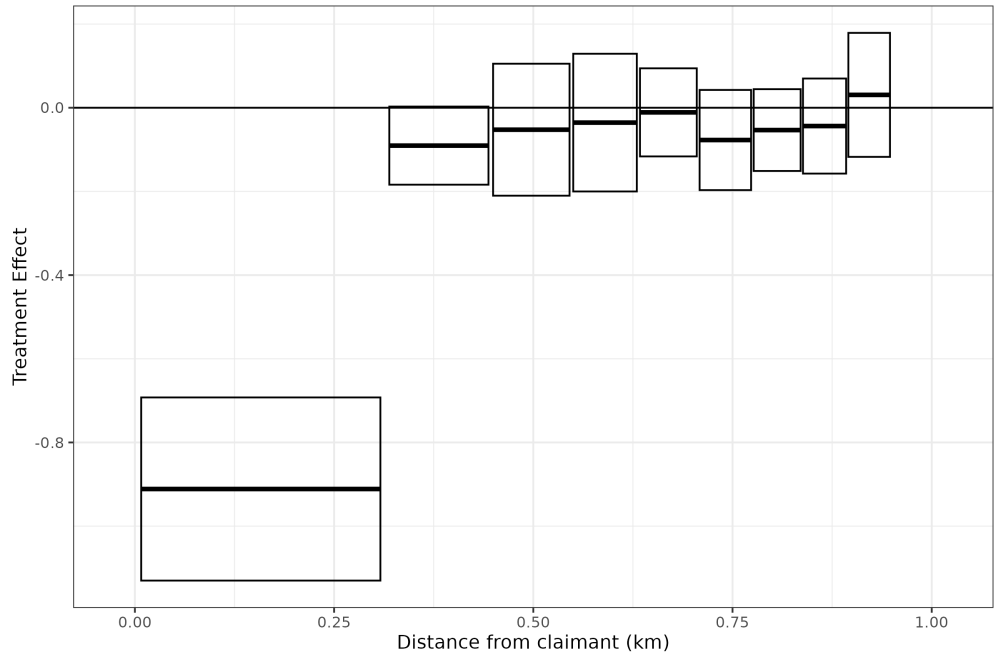


Figure 3: Treatment Effect Dynamics

This figure presents the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for inner ring households (households in postal codes within 100 m of the warranty claim) as compared to outer ring households (households in postal codes between 100 and 200 m of the warranty claim) for each quarter within one year of the claim, using the quarter just before the warranty claim as a baseline. Inner ring households are defined as households in postal codes within 100 m of the warranty claim and outer ring households are defined as households in postal codes within a ring with a width of 100 m outside of the inner ring, with the gap between the inner and outer rings varying in 100 m increments from 0 m to 800 m. The x-axis presents the quarter of the treatment effect relative to the quarter of the corresponding warranty claim. The solid line is the estimate of the treatment effects across the quarters and the range bars depict the 95% confidence interval for the estimates.



(a) Non-Parametric Estimate of Treatment Effect Using Equal Width Rings



(b) Non-Parametric Estimate of Treatment Effect Using Equal Area Rings

Figure 4: Non-Parametric Estimate of Treatment Effect

This figure presents the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for households in postal codes within nine concentric rings around the warranty claim as compared to hoseholds in postal codes within a tenth concentric ring beyond these nine rings (as depicted in Figure 2). In Panel (a), the ten rings are defined to have equal widths, with each of the ten rings encompassing one-tenth of the distance from the claimant postal code centroid to the boundary of a 1 km ring around the claim. In Panel (b), the ten rings are defined to have equal areas, with each of the ten rings encompassing one-tenth of the area of a 1 km ring around the claim. The x-axis presents the distance of the household from the warranty claimant. The thick solid line is the estimate of the treatment effect for each ring, the vertical extents of the rectangle around the thick line depict the 95% confidence interval for each estimate, and the horizontal width of the rectangle depict the width of each of the nine rings.

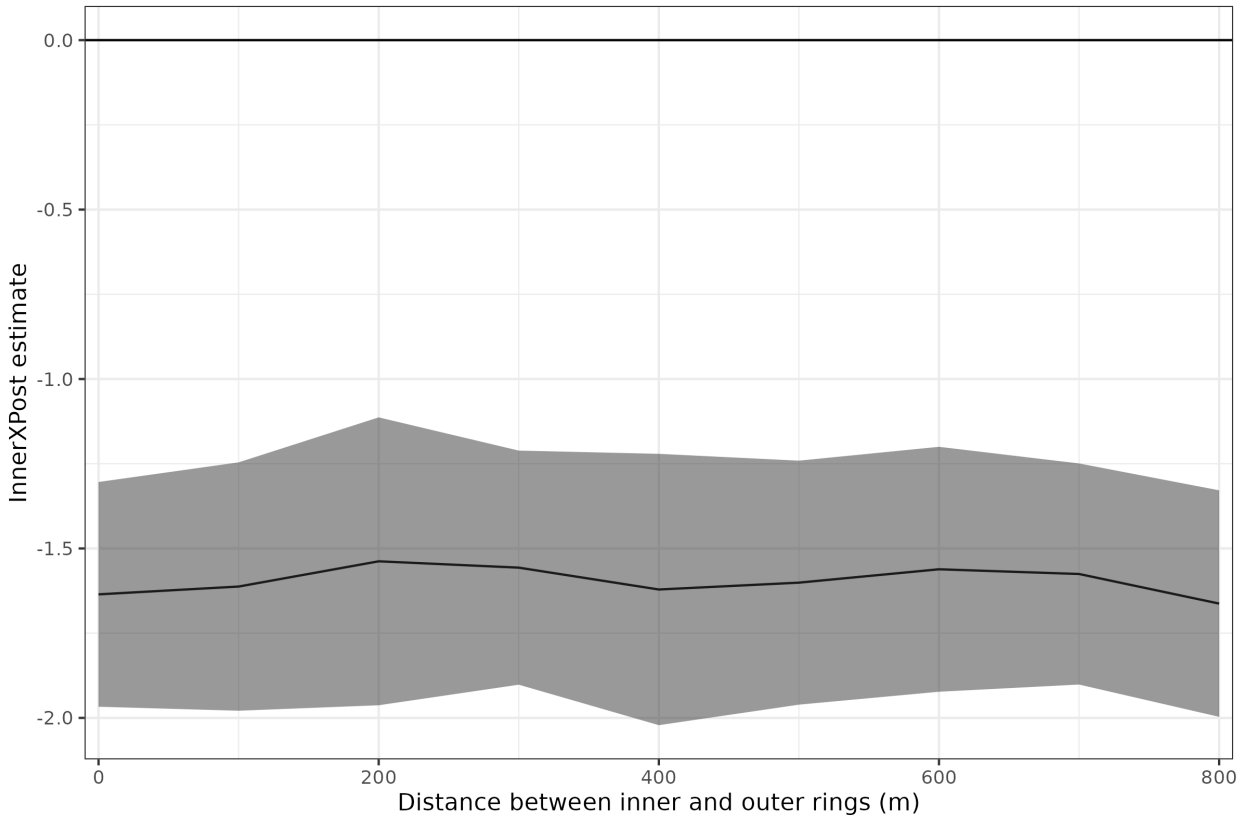


Figure 5: Treatment Effects by "Missing" Distance between Inner and Outer Rings
This figure presents the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for inner ring households as compared to outer ring households. Inner ring households are defined as households in postal codes within 100 m of the warranty claim and outer ring households are defined as households in postal codes within a ring with a width of 100 m outside of the inner ring, with the gap between the inner and outer rings varying in 100 m increments from 0 m to 800 m. The x-axis presents these gaps (in meters) between the inner and outer rings. The solid line is the estimate of the treatment effects across the gap values and the shaded area depicts the 95% confidence interval for the estimates.

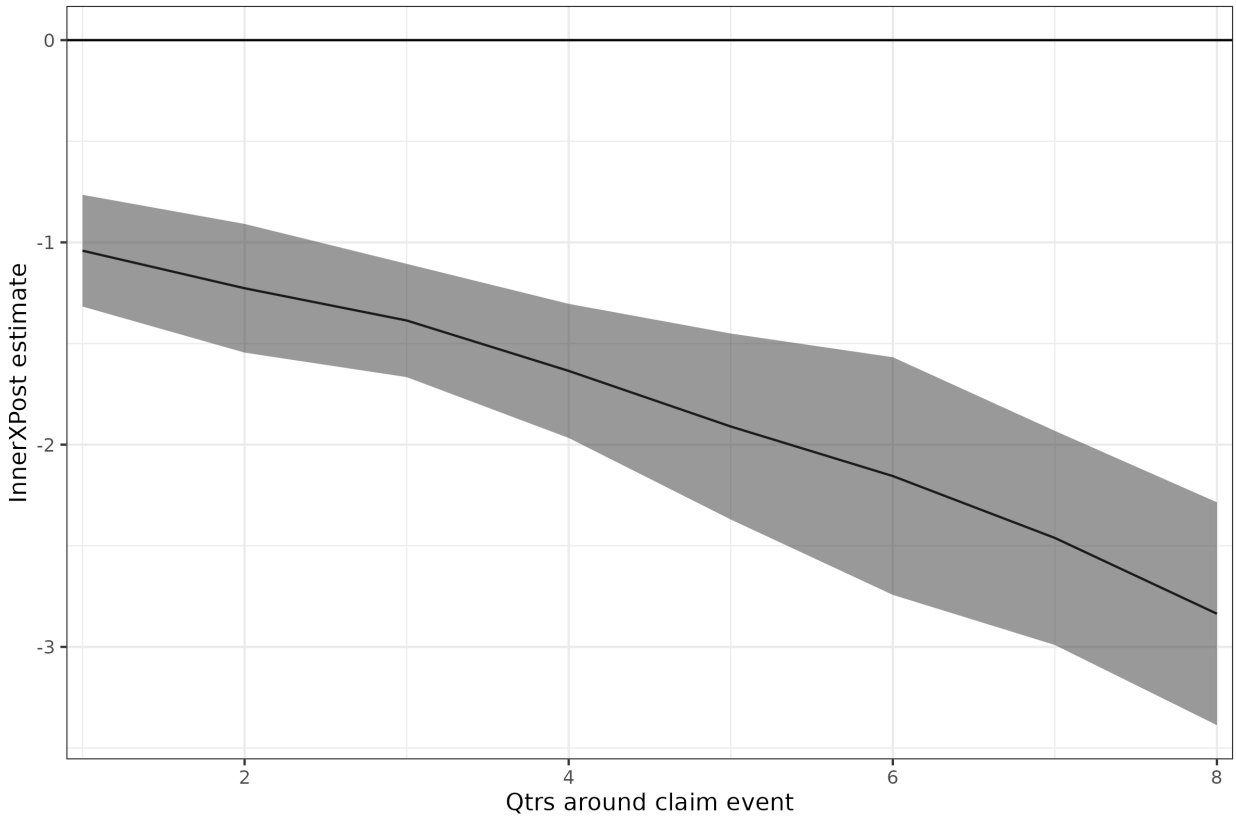


Figure 6: Treatment Effects by Purchase Data Window Length

This figure presents variation in the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for inner ring households (households in postal codes within 100 m of the warranty claim) as compared to outer ring households (households in postal codes between 100 and 200 m of the warranty claim) using purchase data for differing numbers of quarters around warranty claims, which can vary from one to eight quarters. The x-axis presents the number of quarters of purchases used in the analysis. The solid line is the estimate of the treatment effects across the purchase data window lengths and the shaded area depicts the 95% confidence interval for the estimates.

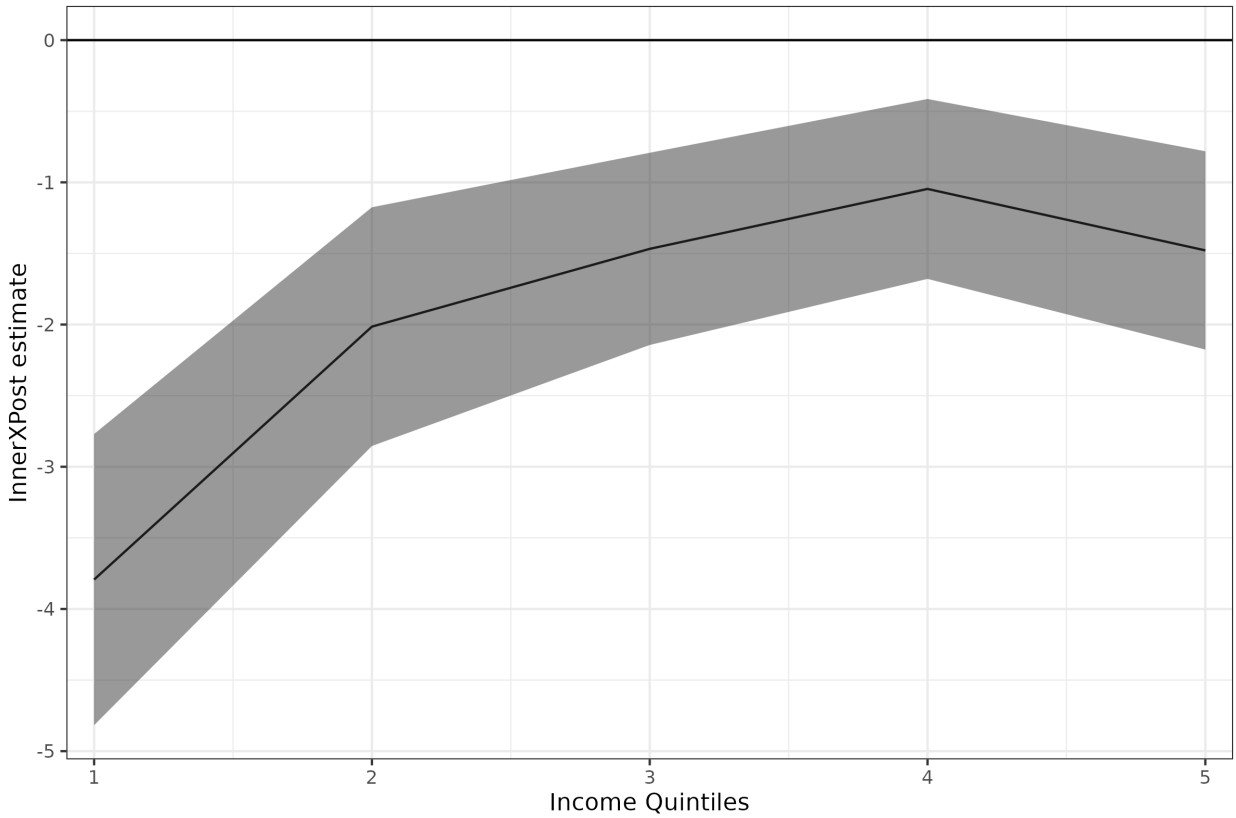


Figure 7: Treatment Effect by Neighborhood Average Income

This figure presents variation in the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for inner ring households (households in postal codes within 100 m of the warranty claim) as compared to outer ring households (households in postal codes between 100 and 200 m of the warranty claim) across the average income quintiles of the claimant household's dissemination area (DA). The x-axis presents the average income quintiles. Quintile 1 is the poorest quintile. The solid line is the estimate of the treatment effects across the income quintiles and the shaded area depicts the 95% confidence interval for the estimates.

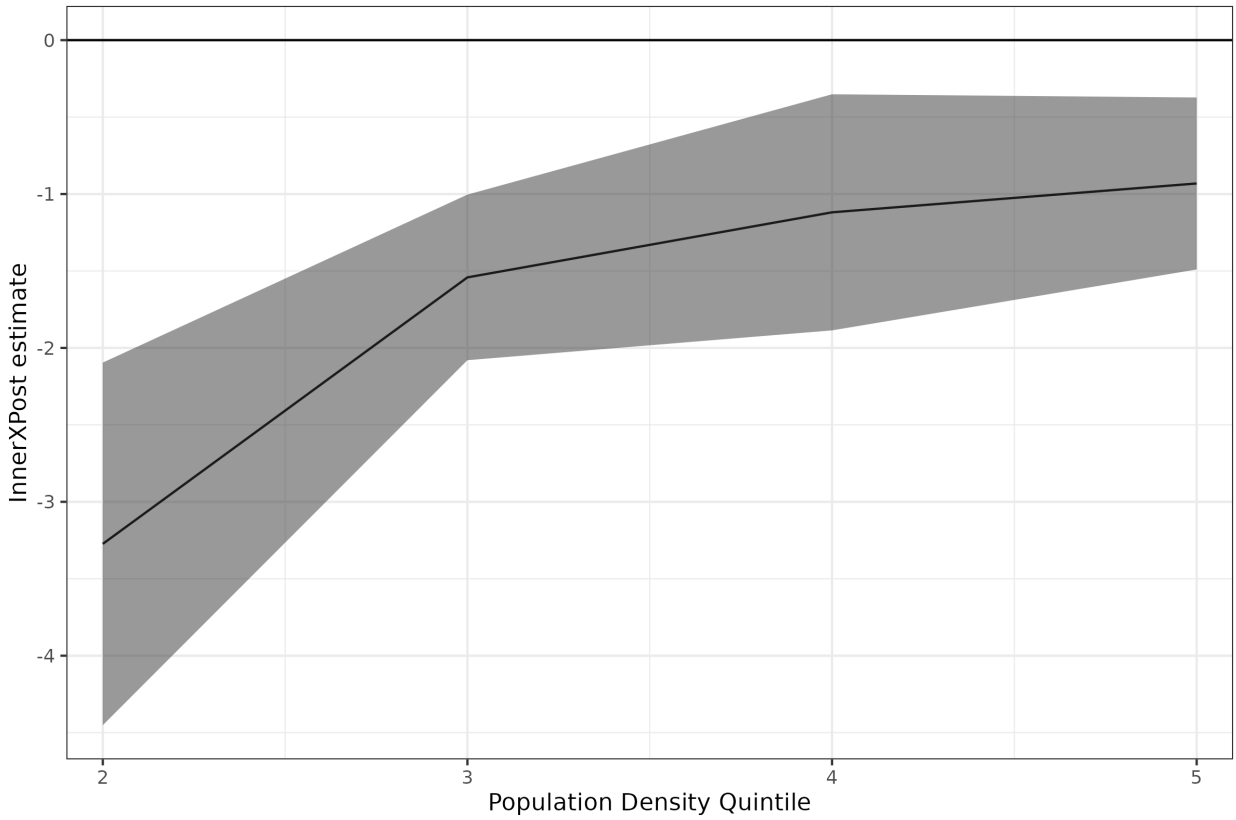


Figure 8: Treatment Effect by Neighborhood Population Density

This figure presents variation in the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for inner ring households (households in postal codes within 100 m of the warranty claim) as compared to outer ring households (households in postal codes between 100 and 200 m of the warranty claim) across the population density quintiles of the claimant household's dissemination area (DA). The x-axis presents the population density quintile of the DA, with the lowest density quintile, which is composed of rural areas, omitted in the figure as it is in analyses. The solid line is the estimate of the treatment effects across the population density quintiles and the shaded area depicts the 95% confidence interval for the estimates.

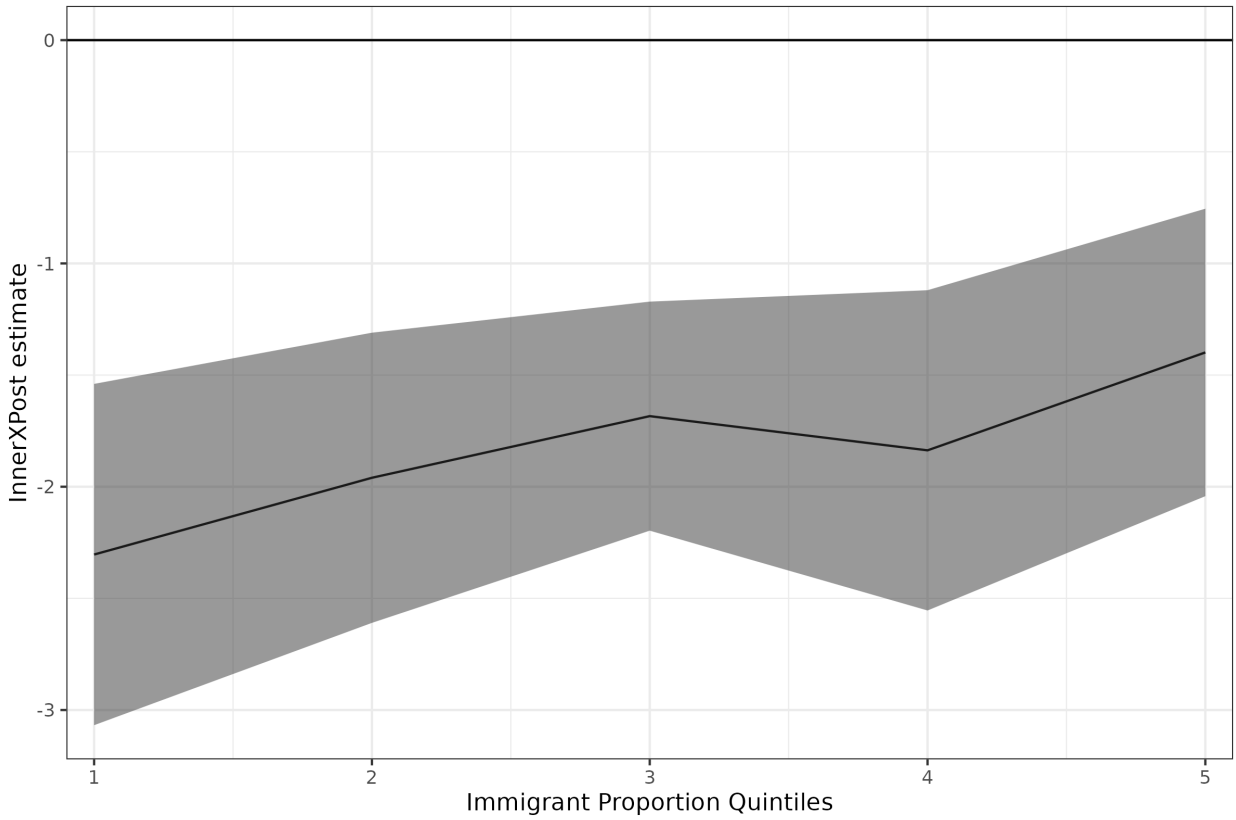


Figure 9: Treatment Effect by Neighborhood Immigrant Proportion

This figure presents variation in the treatment effect of a warranty claim on the purchases of claimed brand goods relative to other brands' goods for inner ring households (households in postal codes within 100 m of the warranty claim) as compared to outer ring households (households in postal codes between 100 and 200 m of the warranty claim) across quintiles of immigrant proportion of population in the claimant household's dissemination area (DA). The x-axis presents the immigrant proportion quintiles. The solid line is the estimate of the treatment effects across the quintiles and the shaded area depicts the 95% confidence interval for the estimates.

Table 1: Summary Statistics

These tables present summary statistics for key variables in this paper. Panel A provides statistics on purchases, Panel B provides statistics on all warranty claims, Panel C lists the top 30 brands by purchases along with purchase market share statistics. And Panel D provides statistics on warranty claims used in our analyses, and Panel D provides demographics for dissemination areas within and outside our data.

(a) Purchases

This panel provides summary statistics for purchases of all consumer durable goods purchased from an anonymous Canadian retail store chain between 1998 and 2009. *Purchase Price* is the price of the product purchased (in Canadian Dollars). *Warranty Purchased* is an indicator variable for the purchase of warranty coverage together with the product by a consumer. *Cost of Warranty* is the price of the purchased warranty coverage (in Canadian Dollars).

	25th %ile	Median	75th %ile	Mean	Std Dev	Obs
Purchase Price	199.97	439.96	788.00	609.27	1765.41	5956402
Warranty Purchased	0.00	0.00	1.00	0.36	0.48	5956402
Cost of Warranty	9.99	69.99	124.99	90.56	95.24	2136208

(b) Top 30 Brands by Purchase Market Share

This panel lists the top 30 manufacturer brands by total purchases of consumer durable goods made between 1998 and 2009 across all branches of an anonymous Canadian retail store chain. For each brand, the panel provides the number of purchases recorded as well as the percent of all purchases.

Manufacturer	Count	Percent
Sony	796,122	14.2
Frigidaire	664,920	11.9
Panasonic	360,781	6.4
Whirlpool	309,678	5.5
LG Electronics	307,441	5.5
Samsung	302,170	5.4
Citizen	223,800	4.0
GE Appliances	211,015	3.8
RCA	206,165	3.7
Maytag	199,213	3.6
Danby	136,839	2.4
Electrolux	122,089	2.2
Inglis	114,149	2.0
Performa	91,319	1.6
Hitachi	87,353	1.6
JVC	82,446	1.5
Sharp	78,693	1.4
Toshiba	75,963	1.4
GTech	75,678	1.3
Prima	61,459	1.1
Brada	56,949	1.0
Stitch	48,566	0.9
Kitchenaid	44,682	0.8
Legend	40,109	0.7
Vitel Malta	38,587	0.7
Moffat	35,610	0.6
Hoover	32,610	0.6
Goldstar	29,751	0.5
Philips	28,951	0.5
Retailer	28,668	0.5

(c) Claims

This panel provides summary statistics for warranty claims on consumer durable goods purchased from an anonymous Canadian retail store chain filed between 1998 and 2009. The summary statistics provided in the four leftmost columns are for warranty claims that are used for analysis as per the procedure detailed in Section 3, whereas the summary statistics provided in the four rightmost columns are for all warranty claims. *Purchase Price* is the price of the product purchased (in Canadian Dollars). *Cost of Claim* is the repairing cost of the product under warranty (in Canadian Dollars). *Days to Claim* is the number of days from the purchase date to the warranty claim date. *Days to Repair* is the number of days spent in repairing the product. *In Analysis Sample* is an indicator variable for whether a claim is used for our analyses.

	Analysis Claims				All Claims			
	Median	Mean	Std Dev	Obs	Median	Mean	Std Dev	Obs
Purchase Price	877.94	1140.21	864.99	1776	769.97	1030.17	816.75	95109
Cost of Claim	128.57	210.67	249.82	1776	118.77	175.63	202.78	95109
Days to Claim	668.00	743.73	417.60	1776	664.00	725.19	423.30	95109
Days to Repair	8.00	18.03	33.18	1776	9.00	19.06	36.53	95109
In Analysis Sample	1.00	1.00	0.00	1776	0.00	0.02	0.14	95109

(d) Dissemination Area Demographics

This panel provides demographics for dissemination areas (a geographic area comprising several postal codes) across Canada. The dissemination areas (DAs) are separated into in-sample and out of sample DAs, where in-Sample DAs are DAs containing at least one warranty claim event. *Population* is the number of people living in a DA. *Pct. Male* is the percentage of the DA population that is male. *Pct. Married* is the percentage of the DA population that is married. *Pct. Immigrants* is the percentage of the DA population that has immigrated to Canada. *Avg. Household Size* is the average size of households in the DA. *Avg. Income* is the average income of adults in the DA (in Canadian Dollars).

	In-sample DAs				Out of sample DAs			
	Median	Mean	Std Dev	Obs	Median	Mean	Std Dev	Obs
Population	557.00	689.88	798.63	1477	516.00	591.39	383.06	51496
Pct. Male	50.61	50.41	2.88	1474	49.26	49.04	3.42	51406
Pct. Married	56.73	53.55	12.14	1474	49.34	47.61	13.52	51406
Pct. Immigrants	7.27	9.50	8.76	1424	11.76	17.85	17.99	50590
Avg. Household Size	2.60	2.61	0.46	1477	2.60	2.58	0.55	51496
Avg. Income	32158.00	34371.32	17959.15	1477	31388.00	33922.92	18089.29	51496

(e) Quintile Medians

This panel provides median values for each quintile of the variables used to study heterogeneities in this paper. *Avg. Income* is the average income of households in analysis sample DAs. *Population Density* is the number of people per square kilometer in dissemination areas (DAs) used in our analysis. *Pct. Immigrants* is the percentage of the population that has immigrated to Canada in analysis sample DAs.

	Median Value for				
	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Avg. Income	21635	27680	32158	37733	48286
Population Density	0.742	9.42	27.3	100	930
Pct. Immigrants	0.806	3.88	7.27	12.1	20.7

Table 2: Dynamic Treatment Effects on Neighbors' Net Purchases of Claimed Brand Products

This table presents the results of geographic dynamic difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code over the course of one quarter. The reported coefficients represent the change in the difference in the dependent variable between inner ring (treated) and outer ring (control) postal codes in the four quarters within one year of a warranty claim, relative to the difference in the quarter before the warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. The twelve columns represent four groups of three regressions where the groups have four different widths for the rings (100 m, 200 m, 300 m, and 400 m) and the three regressions within each group have differing distances between the inner and outer rings (0 m, 100 m, and 200 m). For all regressions, we impose postal code, relative quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	T: 0-100m			T: 0-200m			T: 0-300m			T: 0-400m		
	C: 100-200m	C: 200-300m	C: 300-400m	C: 200-400m	C: 300-500m	C: 400-600m	C: 300-600m	C: 400-700m	C: 500-800m	C: 400-800m	C: 500-900m	C: 600-1000m
Inner x Qtr - 4	0.055	0.045	0.045	0.032	0.019	0.015	0.012	0.014	0.037	0.024	0.022	0.023
	[0.519]	[0.449]	[0.434]	[0.415]	[0.246]	[0.214]	[0.219]	[0.250]	[0.726]	[0.602]	[0.564]	[0.556]
Inner x Qtr - 3	0.011	0.003	0.001	-0.001	0.008	0.038	0.026	0.025	0.030	0.027	0.027	0.011
	[0.114]	[0.029]	[0.014]	[-0.010]	[0.115]	[0.557]	[0.536]	[0.547]	[0.624]	[0.688]	[0.726]	[0.268]
Inner x Qtr - 2	0.034	0.061	0.049	0.046	0.049	0.065	0.045	0.041	0.032	0.027	0.026	0.014
	[0.309]	[0.675]	[0.513]	[0.666]	[0.687]	[0.978]	[0.870]	[0.786]	[0.628]	[0.697]	[0.651]	[0.348]
Inner x Qtr 0	-0.065	-0.089	-0.077	-0.067	-0.039	-0.028	-0.020	-0.025	-0.030	-0.014	-0.019	-0.019
	[-0.851]	[-1.176]	[-0.897]	[-1.049]	[-0.603]	[-0.449]	[-0.381]	[-0.518]	[-0.624]	[-0.380]	[-0.530]	[-0.549]
Inner x Qtr + 1	-0.099	-0.093	-0.095	-0.070	-0.072	-0.076	-0.057	-0.062	-0.060	-0.047	-0.053	-0.058
	[-1.039]	[-0.952]	[-0.914]	[-0.912]	[-0.926]	[-1.031]	[-0.978]	[-1.171]	[-1.156]	[-1.195]	[-1.339]	[-1.439]
Inner x Qtr + 2	-0.152	-0.170	-0.156	-0.126	-0.134	-0.139	-0.100	-0.106	-0.089	-0.074	-0.077	-0.082
	[-1.185]	[-1.300]	[-1.100]	[-1.176]	[-1.223]	[-1.323]	[-1.209]	[-1.355]	[-1.239]	[-1.360]	[-1.461]	[-1.687]
Inner x Qtr + 3	-0.149**	-0.120**	-0.087	-0.066	-0.051	-0.036	-0.019	-0.025	-0.031	-0.024	-0.024	-0.045*
	[-2.551]	[-2.100]	[-1.618]	[-1.619]	[-1.205]	[-0.877]	[-0.625]	[-0.901]	[-1.038]	[-0.986]	[-0.999]	[-1.768]
Inner x Qtr + 4	-0.194**	-0.159*	-0.145	-0.105	-0.106	-0.109*	-0.078	-0.081	-0.067	-0.055	-0.051	-0.068*
	[-2.131]	[-1.941]	[-1.624]	[-1.640]	[-1.617]	[-1.700]	[-1.466]	[-1.664]	[-1.321]	[-1.397]	[-1.327]	[-1.838]
N	31266	34110	36378	54162	57789	60939	83835	87939	92925	119448	126765	131985
R-sq	0.718	0.718	0.718	0.711	0.707	0.707	0.696	0.692	0.691	0.677	0.676	0.673
Postal code FE	X	X	X	X	X	X	X	X	X	X	X	X
Rel qtr FE	X	X	X	X	X	X	X	X	X	X	X	X
Claim ring FE	X	X	X	X	X	X	X	X	X	X	X	X
Claim quarter FE	X	X	X	X	X	X	X	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 3: Non-Parametric Estimation of Treatment Effects (Butts, 2023)

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim based on the non-parametric methodology proposed in Butts (2023). The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the year before or after a warranty claim. The reported coefficients in the nine columns in both subtables represent the change in the difference in the dependent variable between households in postal codes in an inner ring (treated) and postal codes in a fixed outer ring (control) from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. In Panel (a), each of the ten rings in these specifications has an equal width (100 meters), which results in the outer ring containing all postal codes with a centroid within 900 to 1000 meters of the warranty claim event postal code. In Panel (b), each of the ten rings in these specifications has an equal area (0.314 square kilometers), which results in the outer ring containing all postal codes with a centroid within 949 to 1000 meters of the warranty claim event postal code. The column headings in both tables define the extents of the inner (treated) ring for the specification. For all regressions, we impose postal code, post-claim quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

(a) Non-Parametric Estimates of Treatment Effects using Equal Width Rings

	Inner Ring Bounds								
	0-100 m	100-200 m	200-300 m	300-400 m	400-500 m	500-600 m	600-700 m	700-800 m	800-900 m
Inner x After	-1.663***	-0.027	-0.050	-0.125*	-0.106**	-0.042	-0.062	-0.101**	-0.087*
	[-9.751]	[-0.498]	[-0.514]	[-1.901]	[-2.345]	[-0.626]	[-0.932]	[-2.579]	[-1.962]
N	9944	6316	6948	7452	7754	8152	8364	8862	9380
R-sq	0.911	0.729	0.740	0.750	0.696	0.765	0.702	0.814	0.773
Postal code FE	X	X	X	X	X	X	X	X	X
Post-claim FE	X	X	X	X	X	X	X	X	X
Claim ring FE	X	X	X	X	X	X	X	X	X
Claim quarter FE	X	X	X	X	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

(b) Non-Parametric Estimates of Treatment Effects using Equal Area Rings

	Inner Ring Bounds								
	0-316 m	316-447 m	447-548 m	548-632 m	632-707 m	707-775 m	775-837 m	837-894 m	894-949 m
Inner x After	-0.911***	-0.091*	-0.052	-0.036	-0.011	-0.077	-0.053	-0.044	0.031
	[-8.161]	[-1.901]	[-0.652]	[-0.423]	[-0.206]	[-1.263]	[-1.070]	[-0.755]	[0.407]
N	12032	6242	5826	5422	5202	5228	5288	5058	4970
R-sq	0.909	0.718	0.696	0.743	0.739	0.843	0.789	0.763	0.771
Postal code FE	X	X	X	X	X	X	X	X	X
Post-claim FE	X	X	X	X	X	X	X	X	X
Claim ring FE	X	X	X	X	X	X	X	X	X
Claim quarter FE	X	X	X	X	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 4: Treatment Effects by Distance between Inner and Outer Rings

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the year before or after a warranty claim. The reported coefficient represents the change in the difference in the dependent variable between households in inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code. Outer ring postal codes are postal codes in a 100 m ring around the claimant postal code with differing distances from the inner ring for each column. The nine columns vary the distance between inner and outer ring postal codes by 100 m, with the first column placing them next to each other and the ninth column placing them 800 m apart. For all regressions, we impose postal code, post-claim quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	0-100 v 100-200m	0-100 v 200-300m	0-100 v 300-400m	0-100 v 400-500m	0-100 v 500-600m	0-100 v 600-700m	0-100 v 700-800m	0-100 v 800-900m	0-100 v 900-1000m
Inner x After	-1.635***	-1.612***	-1.538***	-1.557***	-1.621***	-1.601***	-1.561***	-1.575***	-1.663***
	[-9.667]	[-8.624]	[-7.097]	[-8.833]	[-7.937]	[-8.716]	[-8.469]	[-9.466]	[-9.751]
N	6948	7580	8084	8386	8784	8996	9494	10012	9944
R-sq	0.911	0.911	0.911	0.910	0.912	0.909	0.912	0.909	0.911
Postal code FE	X	X	X	X	X	X	X	X	X
Post-claim FE	X	X	X	X	X	X	X	X	X
Claim ring FE	X	X	X	X	X	X	X	X	X
Claim quarter FE	X	X	X	X	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 5: Treatment Effects by Purchase Data Window Length

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the year before or after a warranty claim. The reported coefficient represents the change in the difference in the dependent variable between households in inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code and outer ring postal codes are postal codes between 100 and 200 m (0.2 km) of the claimant postal code. The eight columns vary the number of quarters included in the pre- and post-claim windows, with the first column including one quarter before and after the claim and the eighth column including eight quarters before and after the claim. For all regressions, we impose postal code, post-claim quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	1 qtr	2 qtrs	3 qtrs	4 qtrs	5 qtrs	6 qtrs	7 qtrs	8 qtrs
Inner x After	-1.041***	-1.227***	-1.386***	-1.635***	-1.910***	-2.155***	-2.461***	-2.836***
	[-7.387]	[-7.560]	[-9.701]	[-9.667]	[-8.144]	[-7.191]	[-9.114]	[-10.089]
N	6948	6948	6948	6948	6948	6948	6948	6948
R-sq	0.796	0.871	0.903	0.911	0.906	0.901	0.898	0.895
Postal code FE	X	X	X	X	X	X	X	X
Post-claim FE	X	X	X	X	X	X	X	X
Claim ring FE	X	X	X	X	X	X	X	X
Claim quarter FE	X	X	X	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 6: Treatment Effects by Strictness of Analysis Sample Definition

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the two years before or after a warranty claim. The reported coefficient represents the change in the difference in the dependent variable between households in inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code and outer ring postal codes are postal codes between 100 and 200 m (0.2 km) of the claimant postal code. The two columns vary the strictness of the sample of postal codes examined, with the first column being the more strict sample definition (and also the definition used in the paper). The analysis sample in the first column includes neighborhoods where only one warranty claim occurs in the neighborhood over the entire studied time period and where the neighborhood does not overlap with the neighborhoods of any other warranty claims. The analysis sample in the second column includes neighborhoods where no other warranty claim occurs within one year of the focal claim and where the neighborhood does not overlap with the neighborhoods of any other warranty claims. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	No other claims	
	Ever	Within 8 qtrs
Inner x After	-1.635***	-0.183***
	[-9.667]	[-10.055]
N	6948	368520
Claim events	1776	25483
R-sq	0.911	0.843
Postal code FE	X	X
Post-claim FE	X	X
Claim ring FE	X	X
Claim quarter FE	X	X
SE clustering	Ring, qtr	Ring, qtr

Table 7: Comparison of Peer Effect and Alternative Product Quality Measures

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim which include alternative product quality measures as explanatory variables. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the two years before or after a warranty claim. The reported coefficient for *Inner x After* represents the change in the difference in the dependent variable between households in inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code and outer ring postal codes are postal codes between 100 and 200 m (0.2 km) of the claimant postal code. The five columns differ in the additional explanatory variables included in the DID regression. Column (1) has no additional explanatory variables and acts as a benchmark for the other four specifications. Column (2) includes brand claim rate, which is defined as the total number of claims made against claimed brand divided by the total number of sales of that brand's products for the claimed product category in the two years before each warranty claim, as an explanatory variable. Column (3) includes product market share, which is defined as the number of purchases of the brand's products divided by the total number of purchases for the claimed product category in the two years before each warranty claim. Column (4) includes warranty claim process length, which is defined as the number of days taken to repair the good. Column (5) includes all three of these measures of product quality as explanatory variables. For all regressions, we impose inner ring, post-claim quarter, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Inner x After	-1.841***	-1.906***	-1.816***	-1.840***	-1.903***
	[-5.567]	[-5.129]	[-5.416]	[-5.500]	[-5.082]
Claim Rate		22.109*			18.777
		[1.868]			[1.385]
Market Share			-4.502*		-3.938
			[-1.952]		[-1.285]
Days to Repair				0.006	0.004
				[0.890]	[0.646]
N	10983	9997	10844	10983	9997
R-sq	0.042	0.042	0.042	0.042	0.042
Inner ring FE	X	X	X	X	X
Post-claim FE	X	X	X	X	X
Claim quarter FE	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 8: Treatment Effects by Neighborhood Average Income

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the year before or after a warranty claim. The reported coefficient represents the change in the difference in the dependent variable between inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code and outer ring postal codes are postal codes between 100 and 200 m (0.2 km) of the claimant postal code. The five columns separate warranty claims into quintiles based on the average income level of the dissemination area (DA) of the warranty claim, with the first column having the lowest average income and the last column having the highest. For all regressions, we impose postal code, post-claim quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	Median Income Quintiles				
	Bottom quintile	Quintile 2	Quintile 3	Quintile 4	Top quintile
Inner x After	-3.793***	-2.015***	-1.468***	-1.046***	-1.479***
	[-7.262]	[-4.704]	[-4.257]	[-3.243]	[-4.157]
N	925	1494	1350	1582	1494
R-sq	0.948	0.938	0.872	0.952	0.898
Postal code FE	X	X	X	X	X
Post-claim FE	X	X	X	X	X
Claim ring FE	X	X	X	X	X
Claim quarter FE	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 9: Treatment Effects by Neighborhood Population Density

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the year before or after a warranty claim. The reported coefficient represents the change in the difference in the dependent variable between inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code and outer ring postal codes are postal codes between 100 and 200 m (0.2 km) of the claimant postal code. The four columns separate warranty claims into quintiles based on the population density of the dissemination area (DA) of the warranty claim, with the first column showing results for the second lowest population density quintile and the last column for the highest. The lowest density quintile, which is composed of rural areas, is omitted in the table as it is in analyses. For all regressions, we impose postal code, post-claim quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	Population Density Quintiles			
	Quintile 2	Quintile 3	Quintile 4	Top quintile
Inner x After	-3.273***	-1.542***	-1.119***	-0.931***
	[-5.446]	[-5.617]	[-2.859]	[-3.268]
N	608	784	1058	1962
R-sq	0.917	0.975	0.915	0.802
Postal code FE	X	X	X	X
Post-claim FE	X	X	X	X
Claim ring FE	X	X	X	X
Claim quarter FE	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr

Table 10: Treatment Effects by Neighborhood Immigrant Proportion

This table presents the results of geographic classical difference-in-differences (DID) regressions on the change in neighbors' purchasing behavior following a warranty claim. The dependent variable in all specifications is the difference between the number of purchases of the claimed brand and purchases of all other consumer durable good brands in a given postal code in the year before or after a warranty claim. The reported coefficient represents the change in the difference in the dependent variable between inner ring (treated) and outer ring (control) postal codes from before to after a warranty claim, with the corresponding t-statistic reported in brackets below each coefficient. Inner ring postal codes are postal codes within 100 m (0.1 km) of the claimant postal code and outer ring postal codes are postal codes between 100 and 200 m (0.2 km) of the claimant postal code. The five columns separate warranty claims into quintiles based on the immigrant proportion of the population of the dissemination area (DA) of the warranty claim, with the first column having the lowest immigrant proportion and the last column having the highest proportion. For all regressions, we impose postal code, post-claim quarter, warranty claim event, and claim quarter fixed effects. Standard errors are clustered at warranty claim and claim quarter levels. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

	Immigrant Proportion Quintiles				
	Bottom quintile	Quintile 2	Quintile 3	Quintile 4	Top quintile
Inner x After	-2.304***	-1.960***	-1.684***	-1.837***	-1.399***
	[-5.909]	[-5.912]	[-6.436]	[-5.021]	[-4.259]
N	1266	1438	1236	1442	1548
R-sq	0.956	0.951	0.931	0.822	0.831
Postal code FE	X	X	X	X	X
Post-claim FE	X	X	X	X	X
Claim ring FE	X	X	X	X	X
Claim quarter FE	X	X	X	X	X
SE clustering	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr	Ring, qtr