Preliminary version. Please do not quote.

Revisiting Consumption-Housing Wealth Relationship: Role of Peer Effects

Niloy Bose* Koustuv Lahiri† Hector Tzavellas‡

September 15, 2025

Abstract

This paper seeks to understand how peer effects moderate individual consumption responses to unanticipated changes in housing wealth. We show that the Sandy Hook School shooting in December 2012 provided a large negative exogenous shock to the local housing market, leading to a significant fall in individual consumption. On this information, we overlay a network based on geographical proximity among individuals in Sandy Hook and its surrounding areas to learn about the extent to which peer effects explain the observed variation in individual consumption responses. Our results suggest that almost 37% of the observed changes in consumption are attributable to peer-induced spillovers, highlighting the critical role social networks play in shaping the outcomes of economic shocks.

*Department of Economics, Virginia Tech. Email: `nbose@vt.edu`.
†Department of Economics, Virginia Tech. Email: `lkoustuv@vt.edu`.
‡Department of Economics, Virginia Tech. Email: `hectortz@vt.edu`.

1. Introduction

Housing wealth makes up about half of total household net worth and nearly two-thirds of the total wealth of the median household in the U.S., and a strong link between housing wealth and consumption is clear from historical data. The House Price-Driven Great Recession of the 2000s has renewed interest in identifying the Marginal Propensities to Consume (MPC) derived from housing wealth and understanding how factors—such as age demographics and household balance sheets—affect consumption sensitivities to changes in housing wealth. While emphasizing the importance of these factors, we argue that individuals' economic decisions are also shaped by the behaviors of those around them. Such behaviors are well-supported by theory (e.g., Duesenberry (1948); Bramoullé et al. (2009); Blume et al. (2015)) and backed by extensive empirical research on peer effects in the context of education (Calvó-Armengol et al. (2009); De Giorgi et al. (2010); Sacerdote (2001)), microfinance (Banerjee et al. (2013)), retirement decision (Duflo and Saez (2003)), and most importantly consumption (De Giorgi et al. (2020); Agarwal et al. (2021)) among others. If people adjust their consumption not only based on changes in their own housing wealth, but also in response to changes in their peers' spending, then separating the two effects in the data is a step forward in understanding how networks among individuals shape the impact of an economic shock and its propagation. This is what we seek to achieve in this paper.

Meeting our objective involves overcoming several challenges. The first challenge is to develop a clear identification strategy that links housing wealth shocks to consumption expenditure. This is necessary because of the confounding variation and reverse causality in the housing wealth-consumption relationship. Recently, researchers have been creative in identifying nearly exogenous variation in house price data. Several recent studies have used cross-city or regional data and addressed these challenges by employing Saiz (2010) city-level estimates of housing supply elasticities as an instrument[1] for changes in house prices across cities (Mian et al. (2013); Mian and Sufi (2014)). Others have adapted the instrument by interacting supply elasticities with additional variables, such as the long-term interest rate (Aladangady (2017)) and the historical sensitivity of local and house prices to regional cycles (Guren et al. (2021)).

In contrast to the above strategies, we follow Bose and Murshid (2024) and incor-

[1]Though innovative in its rights, Saiz (2010) instrument's use has met criticism owing to the fact that supply constraints are correlated with other city characteristics (Davidoff (2015)), and the relationship between supply constraints and price volatility is much weaker after accounting for observable factors that drive the demand for housing.

porate a natural variation in our data caused by a shock to the local real estate market following the tragic shooting at Sandy Hook Elementary School (SHES) on December 14, 2012, one of the deadliest acts of school violence in recent history. Immediately after the incident, the school was demolished[2], disrupting the availability of schooling services within the Sandy Hook Elementary School attendance zone (SHSEAZ). The literature provides strong evidence (e.g., Black (1999); Figlio and Lucas (2004)) that real estate prices are closely linked to the quality of school services. Based on this evidence, we hypothesize that the absence of local school services caused a negative shock to the SHSEAZ real estate market. We estimate this exogenous shock as plausibly orthogonal to confounding influences and examine how this shock affects individual consumption responses to gauge sensitivity to house price fluctuations.

Our analysis requires detailed data that enables us to accurately link housing wealth shocks to consumption and also helps us understand the network structure among individuals. We meet part of this requirement with the use of proprietary data compiled by one of the three credit-ratings agencies[3]. The data enable us to estimate individual-level pre and post-shooting spending totals with reasonable accuracy. Besides spending information, our data includes other explanatory variables such as income and demographics. Importantly, what sets our analysis apart from much earlier work is the precise level of geocoding in our data, which allows us to geocode individuals and properties to the nine-digit ZIP code (i.e., ZIP +4). With this level of granularity, we are able to construct a meaningful geographical proximity-based network among the individuals and also map wealth shocks into consumption with precision.

Understanding the network structure among a group of individuals and knowing their behaviors are important, but they are not sufficient to estimate the effects of peers. The challenge of linking peer behavior to individual behavior has been difficult because the correlation in outcomes among peers could be caused by endogenous peer selection and/or common shocks. Even without correlated effects, the contextual (exogenous) effects and (endogenous) peer effects are observationally indistinguishable because, in a peer group, everyone's behavior influences others. This classic 'reflection problem' (Manski (1993)) must be resolved before peer effects can be identified within a linear-in-means framework. To do this, we follow Calvó-Armengol et al. (2005) to leverage the network structure and use the deviation of an individual's average peer consumption

[2]The children were moved to a different school district, and uncertainty lingered about reopening a new school inside SHSEAZ. A new school was finally built after four years to resume services.

[3]As a part of our data-sharing agreement, we are barred from disclosing the name of this institution.

from the current global mean of peer averages as an instrument. This instrument aims to capture the unique characteristics of one's peer group that affect peers' outcomes but have little or no direct impact on the individual, especially when we leverage panel data to difference out the group's time-invariant characteristics.

Meeting these challenges positions us well to proceed with the main analysis. First, we consider the possibility that the shock to real estate varied discontinuously at the boundary of the Sandy Hook attendance zone and estimate the exogenous shock to house prices using a difference-in-difference approach that compares SHESAZ real estate to properties in other neighborhoods before and after the shooting. Our findings suggest that the shooting had a significant economic impact on a group of homeowners with homes of four or more bedrooms. For this group, the shock caused by the shooting decreased house values by 7%. This outcome is not surprising, as larger homes are typically occupied by families with children. Next, by combining our housing shock estimate with individual-level spending data, we estimate that credit- and charge-card expenditures fell by about 10 cents for each dollar decline in housing equity. Furthermore, to understand how much of this consumption response is driven by peer effects, we estimate these effects in SHESAZ and nearby areas using pre-shooting consumption data. We find that individual consumption is strongly influenced by the consumption of their peers. According to the IV estimate, a one-dollar increase in the average consumption of the peer group results in a 25-cent increase in individual consumption. Finally, we combine this estimate with the network structure in the SHESAZ area to find that nearly 37% of the observed consumption responses in the data can be attributed to peer influence.

There is an emerging consensus among academics and policymakers that moving beyond a representative agent framework and recognizing the importance of heterogeneity are essential for evaluating the outcomes of an economic shock and for sound policymaking (Kaplan and Violante (2018)). This perspective has been strengthened by experiences from the Great Recession, where households and localities that differ in terms of factors such as balance sheets, credit access, income, and unemployment risk responded differently to housing shocks. Our findings add another significant factor to this list by emphasizing the role of social networks in moderating the effects of economic shocks. Acknowledging its importance is crucial not only for understanding how shocks manifest and spread but also for designing prudent macroeconomic policies.

The remainder of the paper is organized as follows. Section 2 provides institutional background on the Sandy Hook event and Newtown's school attendance zones. Section 3 details our data sources and network construction. Section 4 outlines our econometric framework, including the difference-in-differences and instrumental variable models used to identify the wealth shock and peer effects, respectively. We present our main findings in Section 5 providing evidence for the wealth shock's impact and the significant role of peer networks in moderating individual consuption responses. We assess the stability of these findings in Section 6 through a series of robustness and placebo tests. Finally, Section 7 concludes by discussing the policy implications of our findings and the essential role of social networks in propagating economic shocks.

2. Newtown Public Schools and Attendance Zones

There are four Elementary Schools within the Newtown Public School District (NPSD). These are Hawley, Head O'Meadow, Middle Gate and Sandy Hook. These schools are assigned to serve one of four non-overlapping attendance zones (Figure 1). We divide the town of Newton in Figure 1, into two broad areas - the village of Sandy Hook (cross-hatched area in Figure 1) and the rest of the Newton. The Sandy Hook Elementary School serves the Village of Sandy Hook, but the Sandy Hook Elementary School Attendance Zone (SHESAZ) boundaries do not coincide with those of the village of Sandy Hook. Some parcels inside Sandy Hook village are assigned to Middle Gate Elementary, while some, located near the North West border of the village, are assigned to Hawley Elementary. For our analysis, we'll make repeated reference to three non-overlapping areas that together add up to the town of Newtown. The first is the Sandy Hook Elementary School attendance Zone (SHESAZ) (crossed-hatched blue area in Figure 1) which will act as our primary treatment area. The Greater Sandy Hook area will refer to the part of the Sandy Hook village which falls outside the SHESAZ. In Figure 1, this area is also cross-hatched, but shaded in yellow and green. Finally, the Newton Control area will refer to the part of Newton outside the SHESAZ area. This area is marked in yellow, green, and orange in Figure 1.

Central to the research design is the shooting at the Sandy Hook Elementary school on December 14, 2012, which claimed the lives of 20 students and six faculty members. After the shooting, the Sandy Hook Elementary School (SHES) was closed indefinitely and the school building was demolished. The classes were re-housed on the campus of Chalk Hill Middle School, located in an adjacent county seven miles away. This meant that until the construction of the new school building, SHESAZ residents did not have access to a local elementary schooling option. The new school, which was located on the same site as the old one, was opened in August 2016 nearly four years after the shooting. Based on the existing literature connecting school quality and house prices, we form a prior that the Sandy Hook School shooting may have led to a decline in local housing prices. We uncover the exogenous component of the wealth shock to learn about the change in consumption in response to the changes in housing wealth.

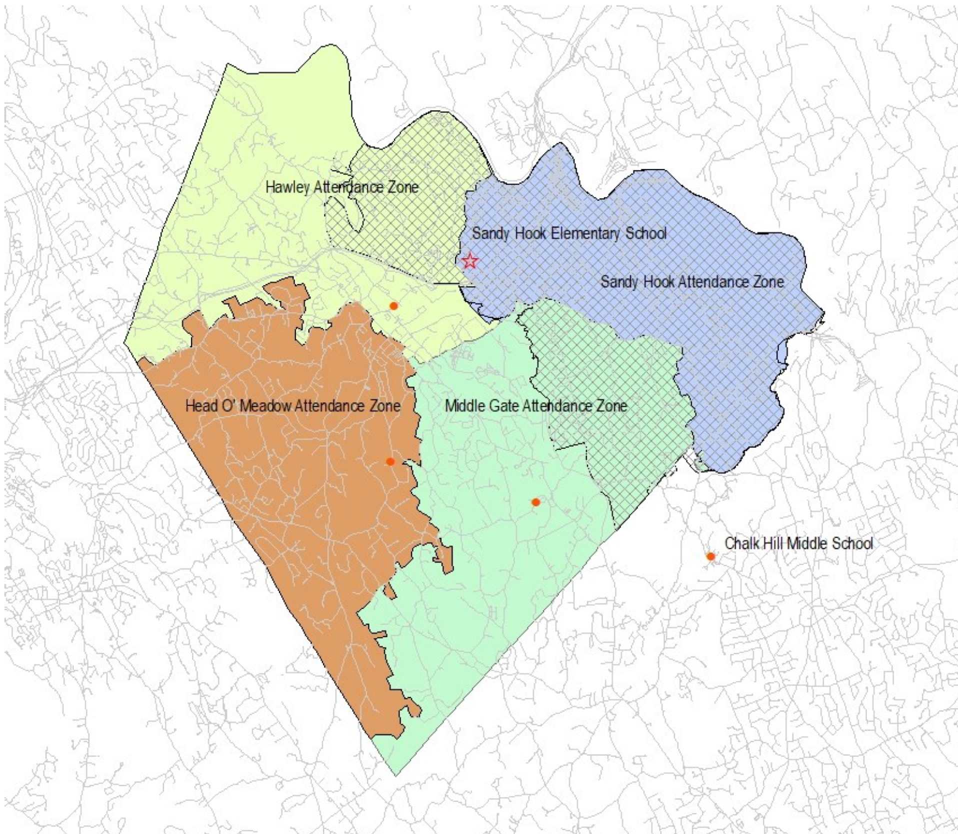

FIGURE 1. Elementary School Attendance Zones in Newtown, CT

3. Data

3.1. Data on Real Estate

We construct a dataset of residential real estate transactions in Newtown, Connecticut, using publicly available records from the Newtown Assessor's Office `https://gis.vgsi.com/newtownct/`. These data, collected via a custom-built Visual Basic web scraper, span the period from January 2003 to July 2016 and include detailed information on transaction prices and physical characteristics of properties sold during this interval.

TABLE 1. Real Estate Transactions: Summary Statistics

	All Areas		SHESAZ		Newtown Control		Gtr. Sandy Hook	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Panel A: All Residential Sales								
Price ($000)	468.4	223.5	414.8	188.9	489.3	232.2	418.1	172.9
Appraisal Value ($000)	394.9	181.7	353.8	159.7	410.9	187.2	364.5	145.7
Square Feet (000s)	2.4	1.1	2.2	1.0	2.5	1.1	2.2	1.0
Lot Size (Acres)	1.8	1.8	1.4	1.3	1.9	1.9	1.6	1.3
Property Grade	82	6	82	5	82	6	81	7
Age at Sale (Decades)	2.9	2.3	3.1	2.5	2.8	2.2	2.3	2.1
Number of Rooms	7.9	2.4	7.7	2.2	8.0	2.5	7.4	2.6
Number of Bedrooms	3.4	0.9	3.3	0.9	3.5	0.9	3.3	1.0
Number of Bathrooms	2.5	0.9	2.3	0.8	2.6	0.9	2.4	0.8
Finished Basement (%)	94	24	97	18	93	25	99	12
Garage (%)	43	49	36	48	46	50	39	49
Central Air (%)	63	48	57	50	65	48	67	47
Hardwood Floors (%)	75	43	72	45	76	43	70	46
Patio/Deck (%)	86	35	85	36	86	35	86	35
Observations	4,203		1,176		3,027		613	
Panel B: Single-Family Homes with 4+ Bedrooms								
Price ($000)	578.4	228.2	534.7	172.7	593.6	242.8	517.7	166.6
Appraisal Value ($000)	497.8	178.6	466.7	134.8	508.6	190.3	458.7	132.2
Square Feet (000s)	3.1	1.0	3.0	0.8	3.2	1.0	2.9	0.8
Lot Size (Acres)	2.3	1.9	2.0	1.5	2.4	2.0	2.1	1.1
Property Grade	85	5	85	4	85	5	85	4
Age at Sale (Decades)	2.0	1.8	1.6	1.6	2.2	1.8	1.9	1.8
Number of Rooms	9.5	1.6	9.4	1.3	9.6	1.7	9.2	1.5
Number of Bedrooms	4.2	0.4	4.1	0.3	4.2	0.4	4.1	0.4
Number of Bathrooms	3.0	0.8	2.8	0.6	3.1	0.9	2.8	0.6
Finished Basement (%)	99	9	99	8	99	10	99	10
Garage (%)	57	50	45	50	61	49	54	50
Central Air (%)	77	42	78	41	76	42	79	41
Hardwood Floors (%)	86	35	86	35	86	35	88	33
Patio/Deck (%)	94	24	94	24	94	24	93	26
Observations	2,218		572		1,646		313	

Notes: Prices and appraisal values are in thousands of dollars. Sales are from January 1, 2003 to July 21, 2016.

A key limitation of the raw dataset is the absence of parcel-level identifiers linking individual properties to their corresponding elementary school attendance zones. To address this, we spatially matched parcels to school catchment areas using geospatial boundary files from the U.S. Census Bureau's Topologically Integrated Geographic Encoding and Referencing (TIGER) database. This process enabled us to accurately classify each transaction as belonging to either the Sandy Hook Elementary School Attendance Zone (SHESAZ), the Greater Sandy Hook area, or the broader Newtown control area. Our final sample consists of 4,203 residential property sales[4]. Table 1 summarizes key features of the dataset, with particular attention to properties featuring four or more bedrooms, which represent approximately half of the local housing stock and form the basis for identifying the housing wealth shock used in our analysis.

3.2. Geographic proximity based Social Network

Our analysis of peer effects requires linking individuals through a social network. To do this, we construct a geographic network based on the spatial coordinates of residential parcels. The process begins with data from the U.S. Census Bureau's Census Geocoder. For each ZIP +4 code in our sample, we collected the latitude and longitude of all constituent parcels, computed the convex hull encompassing them, and assigned the centroid of this polygon as the representative location for that ZIP +4 code.

Using these centroid coordinates as nodes, the peer network is defined through a three-step process:

Step 1: Distance Metric. For any two ZIP +4 centroids i and j, with geographic coordinates $(\text{lat}_i, \text{lon}_i)$ and $(\text{lat}_j, \text{lon}_j)$, we compute the Euclidean distance:

$$d_{ij} = \sqrt{(\text{lat}_i - \text{lat}_j)^2 + (\text{lon}_i - \text{lon}_j)^2}.$$

Given the limited geographic scope of Newtown, this local flat-Earth approximation closely aligns with true geodesic distance.

Step 2: Neighbor Rule. Two centroids i and j are connected if $d_{ij} \leq d^*$, where we set $d^* = 0.004$ decimal degrees (roughly 0.4–0.5 kilometers) to define a neighborhood. However, our main findings are robust to a series of alternative cutoffs. According

[4] In order to improve data reliability and ensure the integrity of our empirical results, we implemented several data-cleaning procedures. We removed duplicate listings, transactions with implausibly low prices (e.g., \$0 or \$1), and non-arm's length transactions such as ownership transfers within families. Additionally, we excluded transactions falling in the bottom and top 1 percent of the sale price distribution to reduce the influence of extreme outliers.

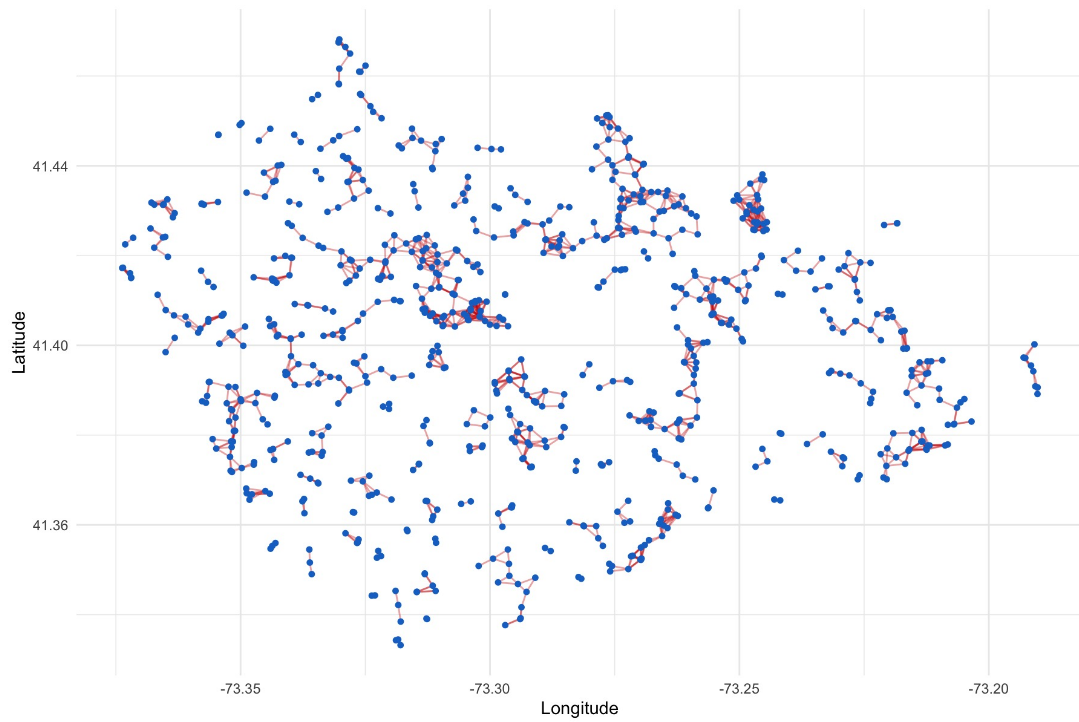

FIGURE 2. Geographic network of ZIP +4 areas in Newtown. Nodes represent ZIP +4 centroids; an edge exists between two centroids if their Euclidean distance is at most $d^* = 0.004$ decimal degrees (approximately 0.45 km).

to this construction, all residents within a ZIP +4 area are mutual peers; if two ZIP +4 centroids are linked, all individuals across both areas are also considered mutual peers.

Step 3: Adjacency Matrix. The binary adjacency matrix $A = [a_{ij}]$ records these network links:

$$a_{ij} = \begin{cases} 1 & \text{if } d_{ij} \le d^* \text{ and } i \neq j, \\ 0 & \text{otherwise.} \end{cases}$$

This adjacency matrix defines the undirected peer network used throughout our analysis.

Figure 2 visualizes the constructed peer network at the ZIP +4 level. In this graph, each node represents a ZIP +4 area, and an edge connects two areas based on a predefined proximity rule. The network translates to individual-level peer relationships as follows: all individuals within a single ZIP +4 node are mutual peers, and if an edge connects two nodes, all residents across both areas become a single, larger peer group. Figure 3 schematically illustrates this translation from geographic links (Panel a) to the

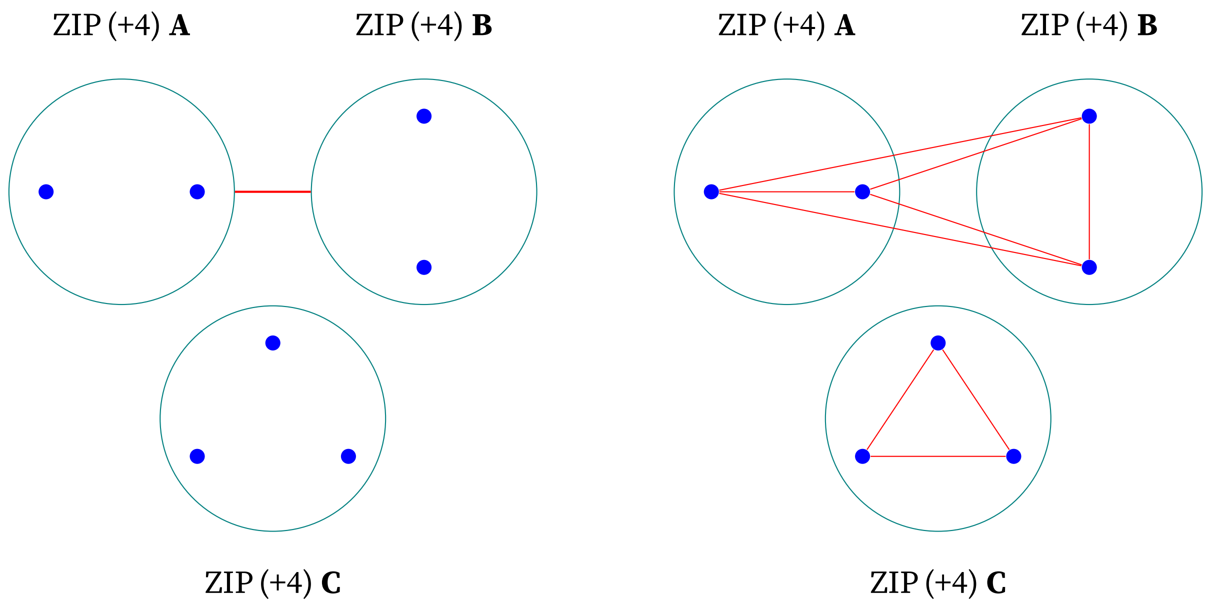

A. ZIP +4 Level Connection B. Resulting Individual-Level Network

FIGURE 3. Schematic of Peer Network Construction. Panel (a) shows a geographic link between two ZIP +4 areas (A and B). Panel (b) shows the resulting individual-level network, where individuals (blue dots) are mutual peers within their own area and also with all individuals in any linked area.

resulting individual-level peer groups (Panel b). An important feature of this pairwise construction is that it generates overlapping, non-exclusive peer groups rather than a simple partition of individuals into distinct clusters. This complex structure is a more realistic representation of social interactions. The descriptive statistics for individuals and their resulting peer groups are summarized in Table 2.

3.3. Tradeline Data

The other data source used in this study consists of detailed financial records collected from monthly tradeline submissions to credit reporting agencies. These data are available for the years 2012 to 2016 covering the pre and post-shooting periods and provide comprehensive information on individuals' revolving credit activity, including credit cards and store charge cards. Notably, the dataset excludes checking account information and other non-credit financial instruments.

Individual consumption expenditures are constructed from these tradeline records by aggregating over accounts and time. For each account k held by individual i in period

TABLE 2. Network Summary Statistics for Newtown Residents

Panel A: Network Structure				
	Mean		SD	
# Direct Neighbors	75.47		37.99	
Panel B: Individual Characteristics				
	Own		Direct Neighbors	
	Mean	SD	Mean	SD
Age (years)	52.77	14.09	50.29	4.75
Family Size	3.71	1.75	3.68	0.49
Married (%)	76.34	42.50	75.19	9.41
Credit Utilization (%)	19.43	23.71	16.61	4.59
Income Insight Score	149.84	76.51	145.78	28.65
Credit Score (VantageScore 3.0)	757.30	71.61	749.46	17.83

Notes: Peer statistics are calculated using direct neighbors within a 0.004 decimal degrees threshold (approximately 0.45 km).

t, we compute monthly expenditures as:

$$X_{ikt} = B_{ikt} - B_{ikt-1} + P_{ikt-1},$$

where B_{ikt} and P_{ikt} denote the account balance and the payment received respectively for individual i on account k at time t. This formulation yields a gross measure of spending that includes payments and the change in outstanding balances. While this measure does not capture total consumption, it represents a consistent and well-defined subset of consumption expenditure.

In addition to spending data, the dataset includes key demographic and financial variables such as income proxies, credit scores (e.g., VantageScore 3.0), credit utilization, and debt-to-income ratios, all of which are observed for the year 2012 at a quarterly frequency. A notable feature of the dataset is the high spatial resolution of individual identifiers: each record is tagged with a 9-digit ZIP code (ZIP +4), which represents a small cluster of residential delivery points—typically around six housing parcels per ZIP +4. While full addresses are omitted to protect anonymity, this granularity enables precise geolocation of each individual and the construction of neighborhood networks based on geographic proximity. Descriptive statistics and variable definitions

TABLE 3. Financial and Demographic Data: Summary Statistics

	All Areas		SHESAZ		Newtown Control		Gtr. Sandy Hook	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Expenditures ($)	4,515	9,341	5,090	9,958	4,389	9,050	5,080	10,149
Balance ($)	3,902	6,395	4,285	6,588	3,843	6,446	3,906	6,100
Credit Limit ($)	35,433	32,580	37,449	33,537	35,302	32,074	36,051	32,307
Credit Utilization (%)	17	22	17	22	16	22	16	21
Debt-to-Income Ratio (%)	15	14	17	14	15	13	16	14
Credit Score	756	73	755	72	757	72	759	70
Income ($000)	75.8	35.6	75.5	34.0	76.9	36.0	74.3	32.1
Age (Years)	53	14	52	13	53	14	53	15
Family Size	3.7	1.8	3.8	1.7	3.7	1.8	3.7	1.8
Married (%)	76	43	78	42	77	42	74	44
Observations	10,822		2,440		7,371		1,451	

Notes: Summary statistics are calculated for individuals with available ZIP code information. "All Areas" refers to Newtown as a whole. Income is reported in thousands of dollars.

are summarized in Table 3.

4. Empirical Strategy

This section outlines our three-stage empirical strategy. First, we detail the difference-in-differences methodology used to identify an exogenous housing wealth shock resulting from the Sandy Hook shooting. Second, we describe our instrumental variable approach to estimate consumption peer effects, addressing classic endogeneity challenges. Finally, we specify the model used to measure how consumption responds to housing wealth shock and how much of this response, observed in the data is due to peer influences.

4.1. Identifying the Housing Wealth Shock

The first stage of our analysis identifies an exogenous variation in housing wealth for the residents of the Sandy Hook Elementary School Attendance Zone (SHESAZ) due to the shooting. Our prior is that the shock to real estate varied discontinuously at the boundary of Sandy Hook attendance zone. With this in mind, we deploy a standard difference-in-difference (DiD) specification as in Bose and Murshid (2024), to identify the causal impact of the school shooting on housing values:

$$\log(p_{irt}) = \alpha + \delta_1 \mathrm{SHESAZ}_i + \delta_2 \mathrm{Post}_t + \beta(\mathrm{SHESAZ}_i \times \mathrm{Post}_t) + X'_{irt}\gamma + \mu_r + \lambda_t + \varepsilon_{irt}, \quad (1)$$

where p_{irt} is the sale price of property i in region r at time t. The variable SHESAZ_i is an indicator equal to 1 if the property i is located within the Sandy Hook Elementary School Attendance Zone (SHESAZ), and 0 otherwise. Post_t is a time dummy that equals 1 for periods following the school shooting incident on Dec 14, 2012. The interaction term $\text{SHESAZ}_i \times \text{Post}_t$ captures the causal effect of the school shooting on home prices in the treatment area. The vector X_{irt} includes a rich set of housing characteristics, such as lot size, square footage, age, and the number of bedrooms and bathrooms. The terms μ_r and λ_t represent the region and time-fixed effects, respectively, to control for unobserved heterogeneity across space and time.

We estimate this regression separately using two control groups. The first specification compares SHESAZ properties to homes in the rest of Newtown (the "Newtown Control"). The second, which we treat as our preferred specification, compares SHESAZ to the "Greater Sandy Hook" area—that is, parcels in Sandy Hook assigned to attendance zones other than SHES. This group is more geographically and demographically similar to SHESAZ, making it a more suitable control.

Our next step is to translate the regression results into a wealth shock measure applicable to all households in our sample, not just those who sold their homes. Relying solely on transacted properties would introduce significant sample selection bias and fail to provide a shock measure for the majority of households. To address this, we use the coefficients estimated in Equation (1) to impute a wealth shock for every 4+ bedroom property in the SHESAZ area, regardless of whether it was sold. For each individual property, we first calculate its predicted price using the estimated coefficients and its specific characteristics, assuming it is in the treated group post-shooting (i.e., with the interaction term's effect included). We then calculate a counterfactual predicted price for the same property by setting the interaction term to zero, simulating the absence of the shooting's impact. The difference between these two imputed values—the factual and the counterfactual—constitutes our property-level estimate of the wealth loss attributable to the shooting. Consequently, all properties outside the SHESAZ area, as well as those inside it with fewer than four bedrooms, are assigned a zero shock by design. We then aggregate these property-level shocks at the ZIP +4 level by computing the average imputed wealth loss. This final measure serves as our spatially disaggregated indicator of the housing wealth shock and is assigned to all individuals residing in the corresponding ZIP +4 area, allowing us to link neighborhood-level shocks to individual-level consumption responses in later sections.

4.2. Estimating Peer Effects

The second stage requires an estimate of consumption peer effects. Identifying peer effects in consumption requires addressing two fundamental challenges: first, clearly defining the relevant network or reference group; and second, accounting for the endogeneity of peers' consumption.

Defining reference groups or social networks in economic research presents inherent challenges, primarily due to limited data availability (De Paula 2017). In an ideal setting, researchers would conduct detailed surveys to map individuals' full social networks—including family members, friends, coworkers, and others—and gather rich socioeconomic information on both sides of each relationship. However, such comprehensive network data are rarely available, with notable exceptions such as the Add Health dataset in the United States and the microfinance client networks studied by Banerjee et al. (2013). As a result, most empirical studies infer network structures based on observable shared characteristics, such as common race, neighborhood, school enrollment, or cohort membership. We study how individuals' consumption responds to nearby peers' spending behavior by using a network defined by geographic proximity. People who live close to each other are likely to observe and influence one another, making spatial proximity a natural definition of peer groups. Leveraging quarterly panel data from before the wealth shock, we estimate how much of an individual's consumption variation is explained by changes in their neighbors' consumption.

Identifying peer effects in economics poses several well-known econometric challenges (Manski (1993); Brock and Durlauf (2001); Moffitt et al. (2001)). In the standard linear-in-means framework, three distinct types of effects must be distinguished: endogenous effects, where an individual's outcome is influenced by the outcomes of their peers (e.g., a person's consumption responds to changes in their neighbors' consumption); contextual effects, where an individual's outcome is shaped by peers' observable characteristics (such as neighbors' income or financial health); and correlated effects, which arise when individuals within a group share unobserved attributes or experience common shocks (such as a local economic downturn or neighborhood-level influences).

We estimate the endogenous peer effects on individual consumption using quarterly panel data from the Newtown population in 2012, ensuring consumption patterns remain unaffected by the subsequent wealth shock.[5] The core empirical model leverages

[5]December spending is treated as predetermined because fewer than fifteen days remain in the quarter after the event.

geographic proximity as defined by the adjacency matrix A constructed in Subsection 3.2, and its row-normalized counterpart G[6], to capture peer influences in consumption. The linear-in-means model in levels, including individual fixed effects, is specified as:

$$C_{it} = \alpha_i + \lambda \cdot GC_{it} + X_{it}\eta + GX_{it}\psi + u_{it}, \tag{2}$$

where C_{it} is the consumption expenditure of individual i in quarter t, and GC_{it} denotes the average consumption of direct neighbors, calculated using the row-normalized adjacency matrix G. Individual characteristics X_{it} include credit utilization, income insight scores, credit scores, and other demographic covariates; GX_{it} represents contextual effects, which control for the influence of neighbors' observable characteristics on individual consumption. The coefficient of primary interest, λ, measures the endogenous peer effect on consumption.

However, identification of λ in Equation (2) presents well-known challenges. Even abstracting from correlated effects, λ cannot be recovered by Ordinary Least Squares (OLS) because of simultaneity between individual and peers' consumption: an individual's spending both influences—and is influenced by—neighbors' spending, violating exogeneity. Consequently, OLS estimates of λ would be biased. We therefore follow Calvó-Armengol et al. (2005) and adopt an instrumental variables (IV) strategy. Our instrument, (Z_{it}) is the deviation of an individual's average peer consumption from the contemporaneous network-wide mean of peer averages, viz.

$$Z_{it} \equiv GC_{it} - \overline{GC_t},$$

where $\overline{GC_t} = \frac{1}{N}\sum_i GC_{it}$. The instrument captures the unique characteristics of one's peer group that affect peers' outcomes which are plausibly exogenous to individual-level consumption decisions.

While, by construction, the instrument is strongly correlated with the endogenous regressor—average peer consumption—thereby satisfying the relevance condition, exclusion could, in principle, fail if Z_{it} correlates with unobserved determinants of i's consumption, u_{it}. In particular, if factors that change the instrument also independently affect C_{it}, the exclusion restriction would be violated, biasing IV estimates (Angrist and Pischke 2009). For instance, shared unobservable traits between i and their neighbors

[6] Specifically, if $A = [a_{ij}]$ is the binary adjacency matrix—where $a_{ij} = 1$ if individuals i and j are connected, and 0 otherwise—then the row-normalized matrix $G = [g_{ij}]$ is defined as $g_{ij} = a_{ij}/\sum_j a_{ij}$. This ensures that GC_{it} represents the average consumption among i's peers in period t.

could induce such a correlation. In our case, this could only happen through correlated effects. We therefore explicitly consider potential sources of correlated effects and implement design choices to neutralize them.

Our primary defense against correlated effects is rooted in the fundamental design of our peer network. By defining connections based on geographic proximity rather than endogenous affiliations (such as workplaces, family, or social clubs), we significantly reduce the scope for individuals to sort into groups based on unobserved characteristics that also drive consumption decisions. For example, networks based on profession could be susceptible to industry-specific income shocks that affect both an individual and their peers. Our geographic design is far less prone to such sources of bias, providing a strong first-line defense against both time-invariant and time-varying correlated effects.

While our network design mitigates peer selection, time-invariant heterogeneity—such as stable preferences, deep-seated financial habits, or persistent neighborhood characteristics shared among neighbors—could still pose a challenge. To address this, we apply a first-difference transformation to all variables in our model, including the instrument (ΔZ_{it}). This transformation purges all time-invariant individual fixed effects (α_i) and any stable, shared unobservables, thereby eliminating a major potential source of bias in peer effects estimation. The resulting first-differenced model is specified as:

$$\Delta C_{it} = \lambda \cdot \Delta GC_{it} + \Delta X_{it}\eta + \Delta GX_{it}\psi + \Delta u_{it}. \tag{3}$$

Finally, we must consider the potential for time-varying correlated effects, as first-differencing does not remove shocks that are both shared among peers and vary over time. In our context, the most salient candidate for such a confounding shock would be the Sandy Hook shooting itself. A specific neighborhood might experience a shared psychological trauma or a localized economic disruption in the event's aftermath, which could influence consumption patterns in a way that is correlated with our instrument. Crucially, our identification strategy is designed to preempt this specific concern. By restricting our estimation to the pre-shock quarters of 2012, we ensure that the shooting itself cannot induce a correlated effect within our sample period. Moreover, the use of a narrow time window within a single, socioeconomically homogeneous town makes the emergence of other systematic, neighborhood-specific economic or credit-market disturbances highly unlikely. As a final empirical validation, we confirm that the instrument is orthogonal to a wide range of individual observable characteristics (including credit scores, income proxies, and credit utilization), which is consistent with the absence of confounding local shocks and supports the validity of the exclusion restriction.

4.3. Measuring the Consumption Response to the Wealth Shock With and Without Peer influence

To fix ideas, consider the following specification:

$$\Delta c_{iA} = \theta \cdot \Delta w_A + X'_{iA}\delta + \varepsilon_{iA}, \tag{4}$$

where Δc_{iA} is the change in consumption for individual i, residing in ZIP +4 A, Δw_A is the exogenous wealth shock facing the individual, and X_{iA} denotes a vector of individual i specific characteristics. The information on Δc_{iA} and X_{iA} can be obtained from the data. For example, in our case, Δc_{iA} is the change in annual card-based consumption for individual i, residing in ZIP +4 A between 2012 and 2016 and X_{iA} denotes a vector of baseline covariates including age, age squared, credit score, credit utilization, family size, and homeownership status. In contrast, Δw_A needs to be estimated using a strategy which is outlined in the Subsection 4.1. In the absence of any influence of peers on consumption behavior, the estimate of θ would represent the sensitivity of consumption to changes in housing wealth - a widely accepted strategy in the consumption-housing wealth literature.

Now consider a modified specification that recognizes the presence of peer influence on individual i's consumption behavior:

$$\Delta c_{iA} = \lambda \cdot \Delta G c_{iA} + \tilde{\theta} \cdot \Delta w_A + X'_{iA}\delta + \varepsilon_{iA}, \tag{5}$$

where G is a row-normalized adjacency matrix that defines the peer network. Here, λ measures the strength of endogenous peer effects—that is, the extent to which an individual's consumption responds to peers' consumption. However, taking specification (5) directly to the entire data spanning both pre and post-shooting periods is problematic due to the fact that the correlated effects may be present in the post-shooting data which could corrupt the estimates of peer effects. Therefore, as outlined in Subsection 4.2, we estimate λ using only pre-shooting consumption data, and use this estimate along with the network structure among the individuals in the SHESAZ area to calculate

$$\widehat{\Delta c}_{iA}^{\text{peer}} = \hat{\lambda} \cdot \Delta G c_{iA}, \tag{6}$$

where $\widehat{\Delta c}_{iA}^{\text{peer}}$ captures the portion of individual consumption change that is attributable to interactions with peers. Finally, we define a change in consumption, net of the peer effect by removing $\widehat{\Delta c}_{iA}^{\text{peer}}$ from consumption series of individuals residing in the treatment area, and regress against the wealth shock to obtain a fair estimate of $\tilde{\theta}$. The differnce between the estimates of θ and $\tilde{\theta}$ informs us the extent to which peer effects moderate consumption sensitivities to wealth shock.

5. Results

We now present the results from implementing our three-stage empirical strategy.

5.1. The Housing Wealth Shock

Table 4 reports the DiD estimates of the impact of the school shooting on 4+ bedroom houses, which are most likely to be affected by changes in school quality due to the presence of school-aged children. The coefficient on the interaction term, β, measures the causal impact of the school shooting on property values within SHESAZ. In our preferred specification with the Greater Sandy Hook control group, we find a statistically significant decline of approximately 7% [≈ exp(−0.071)−1] in post-shooting housing prices for 4+ bedroom homes in SHESAZ. This result establishes a meaningful, exogenous shock to the housing wealth of residents in the affected area.

TABLE 4. Difference-in-Differences Estimates of Housing Price Changes Post-Shooting

	Newtown Control	**Greater Sandy Hook (Preferred)**
	(1)	(2)
SHESAZ × Post-shooting	-0.041^{*}	-0.071^{**}
	(0.022)	(0.031)
Property Characteristics	Yes	Yes
Time Fixed Effects	Yes	Yes
Region Fixed Effects	Yes	Yes
Observations	2,104	861
R^2	0.68	0.61

Note: Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.2. Estimation of Peer Effects

Table 5 presents the results from the first-differenced specification for the peer effects estimation. While the OLS estimate (column 1) is large, it is likely biased upwards. Our preferred IV estimate (column 2) yields a coefficient of 0.241, significant at the 1% level. This implies that a $1.00 increase in neighbors' consumption leads to a statistically and economically significant 24-cent increase in an individual's own consumption.

TABLE 5. Quarterly Peer Effect Estimates with Instrument Diagnostics

	Dependent variable: Δ Consumption	
	OLS	IV
Δ Peer consumption, ΔGC	0.596^{***}	0.241^{***}
	(0.039)	(0.055)
Contextual effects (ΔGX)	Included	
Own covariates (ΔX)	Included	
Cluster level	ZIP +4 code	
Observations	15,426	15,426
R^2	0.073	0.060
Instrument Diagnostics (for column 2)		
First-stage F-statistic (relevance)	20,144.8 (p < 0.001)	
Wu–Hausman test (endogeneity)	$p < 0.001$	

Notes. Robust standard errors clustered at the ZIP +4 level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. The instrument used is the deviation of average peer consumption from the global average across all individuals.

Instrument diagnostics further support our approach. The first-stage F-statistic exceeds 20,000, confirming the instrument's strength, and the Wu–Hausman test ($p <$ 0.001) rejects exogeneity of peer consumption, supporting the need for an instrumental variables approach to address simultaneity and ensure consistent estimation.

5.3. Consumption Response to the Wealth Shock

This section presents the main findings of the paper, implementing the estimation procedure detailed in Section 4.3. The results, shown in Table 6, quantify the sensitivity of consumption to the housing wealth shock and isolate the moderating role of peer effects.

TABLE 6. Wealth Effect on Consumption, 2012–2016

	Dependent variable: Δ Consumption	
	Unadjusted	Peer-adjusted
Housing wealth shock (Δw)	0.107***	0.067***
	(0.021)	(0.021)
Baseline covariates	Included	Included
Observations	2,321	2,321
R^2	0.101	0.083

Notes. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Baseline covariates include credit score, credit utilization, age, age squared, household size, and homeownership status. Peer-adjusted consumption subtracts the estimated network-induced component based on $\hat{\lambda} = 0.241$ from individual-level consumption prior to differencing.

Column 1 of Table 6 reports the estimate of θ from Equation (4). This specification regresses the total change in individual consumption (Δc_{iA}) on the wealth shock. The resulting coefficient, $\hat{\theta} = 0.107$, is statistically significant at the 1% level. As established in our strategy, this estimate represents the overall sensitivity of consumption to changes in housing wealth, a measure that implicitly combines an individual's direct response with any influence propagated through their peer network.

Column 2 presents our estimate of $\tilde{\theta}$ from the modified specification in Equation (5). This coefficient is obtained by first purging the consumption data of the peer influence component. Following our methodology, we use the pre-estimated peer effect parameter, $\hat{\lambda} = 0.241$, to compute $\widehat{\Delta c}_{iA}^{\text{peer}}$, the portion of consumption change attributable to peer interactions, as defined in Equation (6). We then define a change in consumption net of the peer effect by subtracting this component from the total consumption change of each individual. Regressing this net consumption change on the wealth shock yields a fair estimate of $\hat{\tilde{\theta}} = 0.067$, which remains highly significant.

As per our empirical strategy, the difference between the estimates of θ and $\tilde{\theta}$ informs us of the extent to which peer effects moderate consumption sensitivities to the wealth shock. The estimated sensitivity falls from 0.107 to 0.067 once the peer-driven component is removed. This substantial reduction reveals that 37.4% of the total consumption response is attributable to network spillovers [(0.107 – 0.067)/0.107]. This finding demonstrates that social interactions play a powerful moderating role, and any analysis that fails to account for this channel would incorrectly attribute this peer-driven portion to the direct wealth effect alone.

6. Robustness

In this section, we conduct a series of tests to establish the robustness of our core findings. The credibility of our main result—the decomposition of the consumption response—hinges on the consistent estimation of the peer effect parameter, λ. Therefore, our analysis focuses primarily on assessing the stability of this estimate against plausible alternative specifications and threats to identification. We probe the sensitivity of our findings across four dimensions: the specification of the network matrix, the composition of the estimation sample, alternative, more socially meaningful network definitions, and a placebo analysis using counterfactual network structures.

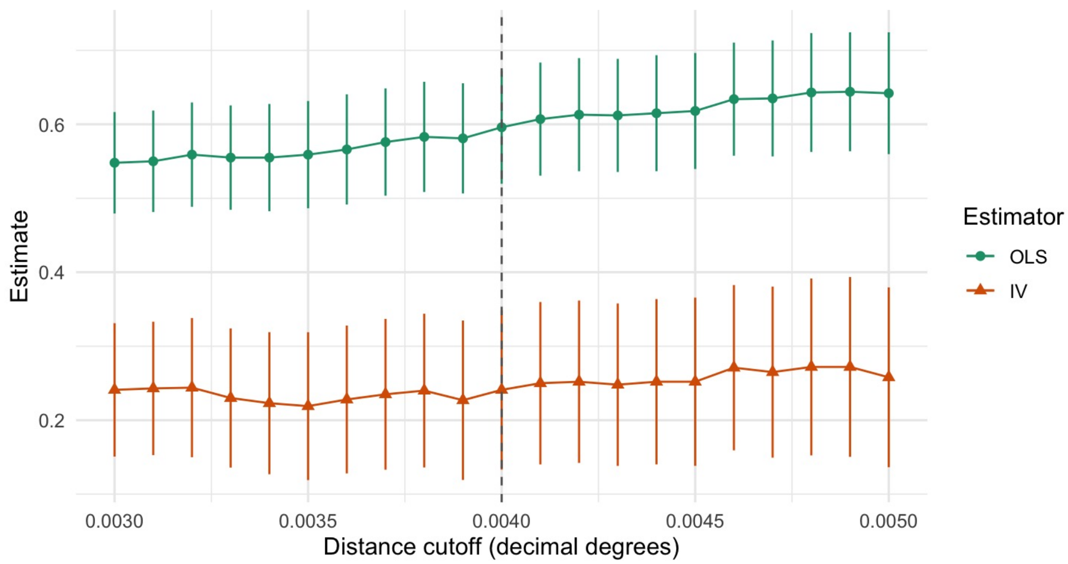

FIGURE 4. Sensitivity to network distance cut-off. Points show point estimates; vertical bars show 95% confidence intervals computed from reported standard errors. The dashed vertical line marks the baseline cut-off (0.004).

First, we examine the sensitivity of our estimate to the specific geographic radius used to define the peer network. Our baseline specification defines neighbors using a 0.004 decimal degree cut-off. In Figure 4, we illustrate the stability of the IV estimate of λ as we vary this threshold from 0.003 to 0.005 decimal degrees (approximately 0.33 to 0.55 kilometers). The estimates remain remarkably stable, ranging from 0.219 to 0.272, and their 95% confidence intervals consistently overlap with that of our baseline estimate of 0.241. This stability suggests that our core finding is not a consequence of an arbitrarily chosen spatial parameter but holds across a plausible range of local neighborhood definitions.

Second, to ensure our results are not an artifact of a specific geographic subsample or driven by unobserved heterogeneity between different areas, we re-estimate our baseline IV model on several distinct partitions of the data. We test the stability of λ across the following subsamples: (i) households within the SHESAZ treatment area only; (ii) households in the local control area (i.e., Newtown excluding SHESAZ); (iii) households within our preferred, wider 'Greater Sandy Hook' control area; and (iv) the full sample of all control households. Across all of these specifications, the IV estimate of λ remains positive, statistically significant, and economically comparable in magnitude to our baseline estimate.

Third, we assess the robustness of our findings to the definition of the peer network itself. While geographic proximity is a natural and observable proxy for social interaction, a potential concern is that it may not be sufficient to capture a truly meaningful social network. To address this, we test whether the peer effect persists when we impose more stringent conditions for what constitutes a peer relationship. We construct a series of alternative networks by by imposing demographic similarity in addition to our baseline geographic proximity criterion. Specifically, we re-estimate our model defining a peer link as two individuals who not only live near each other but also share a key characteristic, such as being in a similar age bracket, having children in the household, or being of the same marital status. These requirements create more restrictive and arguably more socially meaningful peer groups. Although these refined definitions result in sparser networks, which can increase sampling variability, the IV estimate for λ remains positive and statistically significant across all specifications (ranging from 0.168 to 0.280). This demonstrates that our core result is not a product of a simple spatial definition. The peer effect is robust and holds even within these more socially cohesive subgroups, strengthening the interpretation that we are capturing a genuine channel of social influence.

Finally, to validate that the identified peer effect is genuinely rooted in the true spatial structure of the social network, we conduct a placebo analysis. We replace the true adjacency matrix G with randomly generated matrices that preserve certain structural properties of the original network (specifically, we test counterfactuals with the same link density and, separately, the same average node degree). We then re-estimate our baseline IV specification using these counterfactual network structures. Across these placebo scenarios, the IV estimate of λ becomes statistically indistinguishable from zero. This null result provides strong support for our central claim that the actual spatial arrangement of households is the operative mechanism driving the peer effects we document, rather than an artifact of simply including a network structure in the model.

Detailed results and summary statistics from all robustness exercises are provided in Appendix A.

7. Conclusion

This paper investigates how social networks mediate individual consumption responses to economic shocks. By combining a quasi-experimental design with spatial network analysis, we leverage the exogenous variation in housing wealth induced by the Sandy Hook Elementary School shooting to identify a significant and robust peer effect in consumption. Our main contribution is the ability to disentangle the total consumption response into two distinct components: a direct effect from an individual's own wealth shock and an indirect, network-driven spillover effect. We find that these peer spillovers are economically meaningful, shaping the total consumption response as the shock propagates through the local social structure.

The credibility of our findings is grounded in a research design that systematically overcomes two layers of endogeneity. First, the quasi-experimental nature of the housing wealth shock mitigates the reverse causality and omitted variable bias inherent in the consumption-housing wealth relationship, providing an exogenous variation in housing wealth. Second, we address the classic challenges of peer effect estimation—namely, endogenous sorting and the reflection problem (Manski (1993)). Our use of granular ZIP +4 geocoded data allows us to construct a precise proximity-based network, within which we implement an instrumental variable strategy to isolate the endogenous spillover effect. This dual identification strategy, validated by an extensive suite of robustness checks, provides a rigorous foundation for our central claim.

These findings carry significant implications for economic policy. Policies that treat individuals in isolation—without accounting for their social interconnectedness—are likely to miscalculate their ultimate effects. For example, the efficacy of fiscal stimulus, disaster relief funds, or targeted income transfers may depend critically on the social interconnectedness of the recipients. Consequently, understanding local network topography is crucial for designing effective interventions and enhancing the economic resilience of communities by containing the spread of negative shocks.

References

Agarwal, S., Qian, W., and Zou, X. (2021). Thy neighbor's misfortune: Peer effect on consumption. *American Economic Journal: Economic Policy*, 13(2):1–25.

Aladangady, A. (2017). Housing wealth and consumption: evidence from geographically linked microdata. *American Economic Review*, 107(11):3415–3446.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2013). The diffusion of microfinance. *Science*, 341(6144):1236498.

Black, S. E. (1999). Do better schools matter? parental valuation of elementary education. *The quarterly journal of economics*, 114(2):577–599.

Blume, L. E., Brock, W. A., Durlauf, S. N., and Jayaraman, R. (2015). Linear social interactions models. *Journal of Political Economy*, 123(2):444–496.

Bose, N. and Murshid, A. P. (2024). Mpcs of abcs: The housing wealth effect for affluent boomers with credit. *Economic Inquiry*.

Bramoullé, Y., Djebbari, H., and Fortin, B. (2009). Identification of peer effects through social networks. *Journal of econometrics*, 150(1):41–55.

Brock, W. A. and Durlauf, S. N. (2001). Interactions-based models. In *Handbook of econometrics*, volume 5, pages 3297–3380. Elsevier.

Calvó-Armengol, A., Patacchini, E., and Zenou, Y. (2005). Peer effects and social networks in education and crime. Technical report, IUI Working Paper.

Calvó-Armengol, A., Patacchini, E., and Zenou, Y. (2009). Peer effects and social networks in education. *The review of economic studies*, 76(4):1239–1267.

Davidoff, T. (2015). Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors. *Available at SSRN 2400833*.

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

De Giorgi, G., Pellizzari, M., and Redaelli, S. (2010). Identification of social interactions through partially overlapping peer groups. *American Economic Journal: Applied Economics*, 2(2):241–275.

De Paula, A. (2017). Econometrics of network models. In *Advances in economics and econometrics: Theory and applications, eleventh world congress*, pages 268–323. Cambridge University Press Cambridge.

Duesenberry, J. S. (1948). *The consumption function: A study of relations between income and consumer expenditures*. University of Michigan.

Duflo, E. and Saez, E. (2003). 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.

Figlio, D. N. and Lucas, M. E. (2004). What's in a grade? school report cards and the housing market. *American economic review*, 94(3):591–604.

Guren, A. M., McKay, A., Nakamura, E., and Steinsson, J. (2021). Housing wealth effects: The long view. *The Review of Economic Studies*, 88(2):669–707.

Kaplan, G. and Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, 32(3):167–194.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542.

Mian, A., Rao, K., and Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726.

Mian, A. and Sufi, A. (2014). House price gains and us household spending from 2002 to 2006. Technical report, National Bureau of Economic Research.

Moffitt, R. A. et al. (2001). Policy interventions, low-level equilibria, and social interactions. *Social dynamics*, 4(45-82):6–17.

Sacerdote, B. (2001). Peer effects with random assignment: Results for dartmouth roommates. *The Quarterly journal of economics*, 116(2):681–704.

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296.

Appendix A. Robustness Checks

This section reports robustness exercises corresponding to the peer effects estimates in Subsection 5.2. We consider three classes of specification checks: (i) varying the geographic distance threshold used to define peer networks; (ii) re-estimating the model on different regional subsamples; (iii) restricting peer links to demographically similar individuals.

A.1. Distance Threshold Sensitivity

Table A1 reports the OLS and IV estimates for the peer effect parameter λ when the geographic distance cut-off used to define the peer network is varied. The IV estimate remains stable and statistically significant across all tested thresholds. This result is shown in Figure 4.

TABLE A1. Sensitivity to Neighborhood Distance Cut-off

Cut-off	**OLS**	**IV**	**Cut-off**	**OLS**	**IV**
0.0030	0.548^{***}	0.241^{***}	0.0041	0.607^{***}	0.250^{***}
0.0031	0.550^{***}	0.243^{***}	0.0042	0.613^{***}	0.252^{***}
0.0032	0.559^{***}	0.244^{***}	0.0043	0.612^{***}	0.248^{***}
0.0033	0.555^{***}	0.230^{***}	0.0044	0.615^{***}	0.252^{***}
0.0034	0.555^{***}	0.223^{***}	0.0045	0.618^{***}	0.252^{***}
0.0035	0.559^{***}	0.219^{***}	0.0046	0.634^{***}	0.271^{***}
0.0036	0.566^{***}	0.228^{***}	0.0047	0.635^{***}	0.265^{***}
0.0037	0.576^{***}	0.235^{***}	0.0048	0.643^{***}	0.272^{***}
0.0038	0.583^{***}	0.240^{***}	0.0049	0.644^{***}	0.272^{***}
0.0039	0.581^{***}	0.227^{***}	0.0050	0.642^{***}	0.258^{***}
0.0040	0.596^{***}	0.241^{***}			

A.2. Regional Subsamples

Table A2 presents the results from re-estimating our baseline model on distinct geographic subsamples. The peer effect estimate is positive and significant within both the treatment and various control areas, demonstrating the result is not driven by a specific region.

TABLE A2. Estimates by Geographic Subsample

Region	Duration	OLS	IV
Newtown	4 Qtrs	0.596^{***}	0.241^{***}
SHESAZ	4 Qtrs	0.697^{***}	0.382^{***}
Sandy Hook (SHESAZ + Gr SH)	4 Qtrs	0.609^{***}	0.237^{***}
Newtown – SHESAZ	12 Qtrs	0.587^{***}	0.196^{***}
Newtown – SandyHook	12 Qtrs	0.600^{***}	0.234^{***}

A.3. Alternative Network Definitions

Table A3 shows the robustness of our results to more restrictive network definitions that require peers to share demographic characteristics in addition to geographic proximity. The IV estimate for λ remains positive and significant even in these more socially cohesive networks.

TABLE A3. Robustness to Enriched Peer Definitions

Criteria to be a Neighbor	Observations	OLS	IV
Distance $\leq$ Cut-off	15,426	0.596^{***}	0.241^{***}
Distance + Age within 10 years	15,066	0.467^{***}	0.247^{***}
Distance + At least one child	5,814	0.495^{***}	0.280^{***}
Distance + Both married	11,856	0.538^{***}	0.206^{***}
Distance + Same education level	15,020	0.309^{***}	0.168^{***}
Distance + Same occupation	13,822	0.337^{***}	0.179^{***}

Appendix B. Full Regression Output

This section reports the full coefficient estimates for the regression models presented in the main text. Robust standard errors clustered at the ZIP +4 level are shown in parentheses.

B.1. Peer Effects Estimation (Table 5)

Table A4 presents the complete regression output for our main peer effects specification from Table 5. It includes the full set of control variables and their corresponding coefficient estimates for both the OLS and IV models.

TABLE A4. Full Results: Peer Effects Regression (Quarterly Panel, First Differences)

	OLS	**IV**
Peer consumption change (ΔGC)	0.596***	0.241***
	(0.032)	(0.055)
Income Insight Score	2.945***	4.345***
	(0.786)	(1.344)
Avg. Peer Income Insight Score	14.027***	19.488***
	(4.112)	(7.222)
Lagged Credit Utilization	−33.187***	−42.080***
	(2.145)	(3.816)
Avg. Peer Credit Utilization	92.608***	112.187***
	(10.003)	(18.477)
Lagged Credit Score (VantageScore 3.0)	6.109***	8.089***
	(0.898)	(1.547)
Avg. Peer Credit Score	2.001	2.738
	(4.101)	(7.147)
Observations	15,426	15,426
R-squared	0.072	0.060

Note: Robust standard errors clustered at ZIP +4 level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

B.2. Wealth Effect on Consumption (Table 6)

Table A5 expands on the results from Table 6 by providing the full regression output for the estimation of the wealth effect on consumption. The two columns correspond to the unadjusted model and the peer-adjusted model, illustrating how the estimated coefficient on the housing wealth shock changes after controlling for peer consumption.

TABLE A5. Full Results: Wealth Effect on Consumption (Annual Change, 2012–2016)

	Unadjusted	**Peer-adjusted**
Housing Wealth Shock (Δw)	0.107***	0.067***
	(0.021)	(0.021)
Credit Score (VantageScore 3.0)	−41.138***	−40.743***
	(4.408)	(4.397)
Age	−227.063**	−254.421***
	(88.813)	(88.356)
Age Squared	2.727***	2.958***
	(0.863)	(0.859)
Income Insight Score	−25.973***	−23.517***
	(4.995)	(4.928)
Credit Utilization	−55.146***	−53.488***
	(12.655)	(12.703)
Married Male Indicator	920.052*	983.903*
	(552.627)	(547.455)
Homeowner Indicator	718.669	863.297
	(730.762)	(721.368)
Observations	2,321	2,321
R-squared	0.101	0.083

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. This table presents OLS estimates of the impact of the housing wealth shock on annual changes in consumption. Column 1 shows the unadjusted effect. Column 2 includes peer consumption as an additional control. Robust standard errors, clustered at the ZIP +4 level, are reported in parentheses.