Appendix to: Social Learning in the COVID-19 Pandemic: Community Establishments' Closure Decisions Follow Those of Nearby Chain Establishments

Mathijs de Vaan, Saqib Mumtaz, Abhishek Nagaraj, and Sameer B. Srivastava

Haas School of Business, University of California, Berkeley *

*Corresponding Author: Mathijs de Vaan (mdevaan@haas.berkeley.edu). Other authors: saqib@berkeley.edu, nagaraj@berkeley.edu, and sameersriv@berkeley.edu. We thank Adit Jain for excellent research assistance. We thank three anonymous reviewers, the associate editor and department editor of Management Science for their valuable feedback and suggestions. We also thank Paul Gouvard and András Tilcsik for feedback on prior drafts of the paper. The usual disclaimer applies.

Appendix A: Data Description

A1: Data Sources

Safegraph Mobility Data

Our data on visits to establishments come from *SafeGraph*, which compiles location trace data that are obtained by tracking GPS pings from about 45 million mobile devices across the United States. SafeGraph follows a multistep process to ensure privacy and reliability.[1] Although the devices tracked by SafeGraph are not randomly drawn from the population of devices in the United States, a series of analyses conducted by SafeGraph show that sampling bias is minimal and is unlikely to drive the results presented in this paper.[2] We use two data sets provided by SafeGraph: *Places* and *Social Distancing Metrics*.

Places. SafeGraph Places is a data set that builds on GPS pings from mobile devices and aggregates visitor information for over four million places of interest. The Places data include visit counts on a daily basis, visitor home census block group (CBGs), and other brand establishments that are also visited by those who visit a given place of interest. Over 3,400 brands are represented in the data, and these brands vary in size and scope. On the one end of the spectrum are brands that operate at a national scale. At the other end are smaller brands with only a few establishments and operations confined to a single region. For each establishment, geographic coordinates, CBG, brand with which an establishment is associated (if any), and state are also available. The complete schema of the Places data set can be found at this url: https://docs.safegraph.com/docs/places-schema#section-patterns.

Social Distancing Metrics. Using the same mobile device data, SafeGraph calculates social distancing metrics by identifying people's homes or the place where they most frequently spend the night using a 6 week rolling window. Based on this location, SafeGraph then determines whether a device spends the day at the nighttime location. Examples of social distancing variables include: number of devices completely at home, number of devices partly at work, and total number of

[1]Further details about their methodology and protocols can be found at their official blog post.
[2]See this blog post for a detailed analysis of possible bias in the data.

devices. These metrics are available on a daily basis at the CBG level. Further details of the social distancing metrics can be found here: https://docs.safegraph.com/docs/social-distancing-metrics.

A2: Defining a National Brand

Our instrumental variable strategy requires a robust definition of what constitutes a *national brand*. The starting point for identifying national brands is the *brands* indicator in SafeGraph's Places data. For each establishment that is owned or associated with (i.e., franchises of) a larger organization, this indicator is set to 1. Figure A1 (left panel) builds on these data, as well as location data, and shows the distribution of the number of states that brands operate in. The right panel shows the number of establishments within each brand. Figure A2 depicts a scatter plot of these two parameters across all brands.

Figure A1: Distribution of brand across states and stores

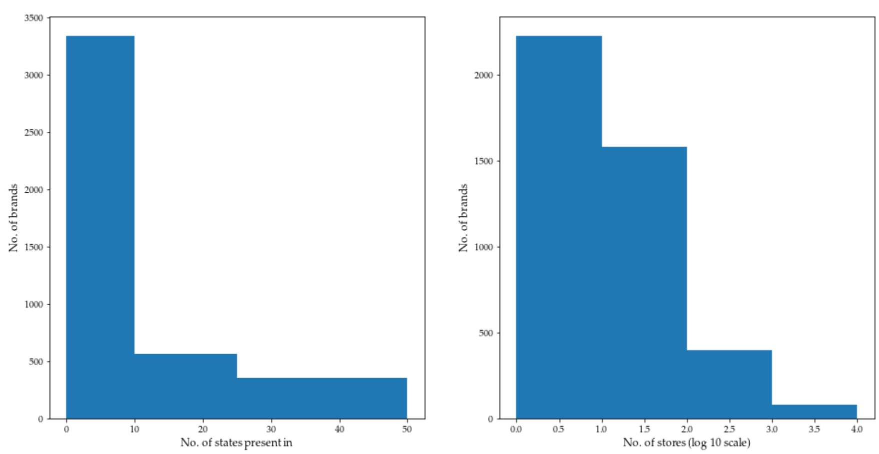

Note: The left panel shows the distribution of number of states that brands operate in, the right panel shows the number of affiliated or owned establishments.

Both figures show that there are a substantial number of brands active in one or only a few states. To ensure that centralized closure decisions were not influenced by idiosyncratic features of a given geography, we defined *national brands* as those with establishments in a large number of states and regions but did so in a way that obtained a sufficiently large sample. We did so by varying the thresholds for both the *number of states* a brand is active (R1) in and the *number of establishments* a brand operates (R2).

Figure A2: Number of stores and states by brand

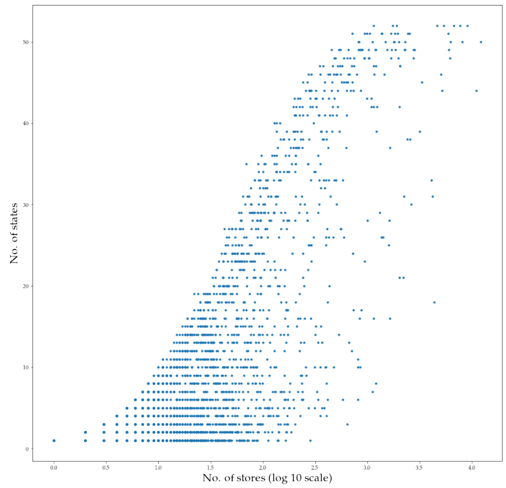

Note: The scatter plot shows that, although the number of affiliated or owned establishments and the number of states a brand is active in are correlated, some brands include a large number of establishments that are concentrated in a small number of states.

Table A1 shows how the number of establishments varies as a function of changing the threshold for the minimum number of states a brand is active in, while Table A2 shows how the number of establishments changes when we change the threshold for the minimum number of stores a brand operates. Inspecting these two tables, we defined brands as *national brands* when they were active in at least 25 states and operated 50 or more establishments.

Table A1: Number of establishments and brands making the cutoff as a function of R1

No. of states cutoff (R1)	No. of establishments	No. of brands
50	75943	26
42	157162	120
25	198551	327
13	224580	665
1	261364	3407

Further, we deemed a chain establishment to be in the same industry as a community establishment if both shared the same 3-digit NAICS code. Since NAICS codes are nested (e.g., 4-digit NAICS codes represent different subgroups of 3-digit NAICS codes), this allowed us to compare establishments in broadly similar industries, while at the same time making sure that a sufficiently

Table A2: Number of establishments and brands making the cutoff as a function of R2

No. Of stores cutoff (R2)	No. of establishments	No. of brands
1,000	122,422	48
500	157,795	98
100	214,796	345
50	231,930	585
25	243,521	912
1	260,951	2,993

high number of community establishments matched with another chain establishment in their zip code.

A3: Determining when an establishment is closed or open

The SafeGraph data do not include an indicator for whether an establishment was open or closed. We therefore developed an algorithm that, for each establishment date, classifies an establishment as being open or closed. We performed two checks to validate this approach, and we developed two simple alternatives that produced nearly identical results.

Our starting point for identifying whether an establishment is open or closed was examining the ratio of visitors on a given day and the average number of visitors on the same weekday in February (before the COVID-19 pandemic led people to begin sheltering-in-place). We call this ratio R_o. Simply using this naive ratio by setting a threshold to label an establishment as closed reveals a set of challenges:

- Daily visit data are variable and can be noisy. Simply setting a threshold for R_o to determine whether an establishment is open or closed resulted in a considerable number of establishments frequently flipping from open to closed and vice versa. To mitigate this issue, we experimented with using a rolling average over N days to refine R_o.
- The threshold set for R_o can be misleading for both large and small establishments. Assume, for example, that we set the threshold to 0.3. A small establishment with five customers on a day in February and one customer on the same day in March would be labeled as closed.

Likewise, an establishment with 500 customers on a day in February and 100 customers on the same day in March, would be labeled as closed. It is likely, however, that the small establishment had closed (and the one visitor is the owner or an employee), while the larger establishment is still open. To address this issue, we experimented with upper and lower limits of the absolute number of visitors in March, L_h and L_l.

- A drop in customers may be demand driven. For example, rather than the owner of an establishment deciding to close their establishment, customers may have started taking precautions, leading them to reduce their foot traffic. To address this issue, we considered the slope of R_o over time. We reasoned that, when establishments closed down, we would observe a more discontinuous change in foot traffic relative to the case where the decline was the result of changing customer preferences. Although our models control for social distancing, a carefully defined cutoff for the grade of the slope may still improve the validity of our outcome measure. To compute the slope and determine a value to adjust our basic algorithm, we need to set two parameters: M, the number of days over which the slope is computed, and G, the grade of this slope.

We experimented with a variety of different values for these inputs and validated the results in two ways. First, SafeGraph provides a small sample of hand-coded closing decisions for certain establishments. They developed a machine learning algorithm to predict the closing status for a slightly larger sample of businesses in a handful of industries. See the SafeGraph website for more details. We took this small sample and compared their classification to that obtained from our strategy using multiple combinations of the parameters described above. We found that setting parameters values equal to $N = 4$, $L_h = 50$, $L_l = 7$, $M = 3$, $G = -0.12$, and $R_o = 0.2$ fit the SafeGraph data most accurately. In other words, we first took a rolling average of the number of visitors over four days and decided that a place was closed if its ratio over a similar day in February was lower than 0.2. We labeled as establishment as being open if it had more than 50 visitors and closed if it had less than 7 visitors. We found a 77.84 % overall overlap between our metric and SafeGraph's indicator. Note that the SafeGraph data should not be considered the "ground truth" on closure. In fact, they use a procedure that is similar to ours in some respects but also makes somewhat different assumptions. Thus, it is not surprising that our determinations of the closure status of an establishment do not exactly match theirs.

Second, we visually inspected a random sample of businesses and plotted visits, as well as our metric of open versus closed. We did this for 256 non-essential businesses (using 4x4 grids). A sample of these visualizations (Figure A3) suggest that our strategy worked quite well. Here a vertical line indicates that an establishment is deemed to be closed according to our algorithm.

Figure A3: Visualising our metric's performance on 16 randomly chosen brands from non essential industries.

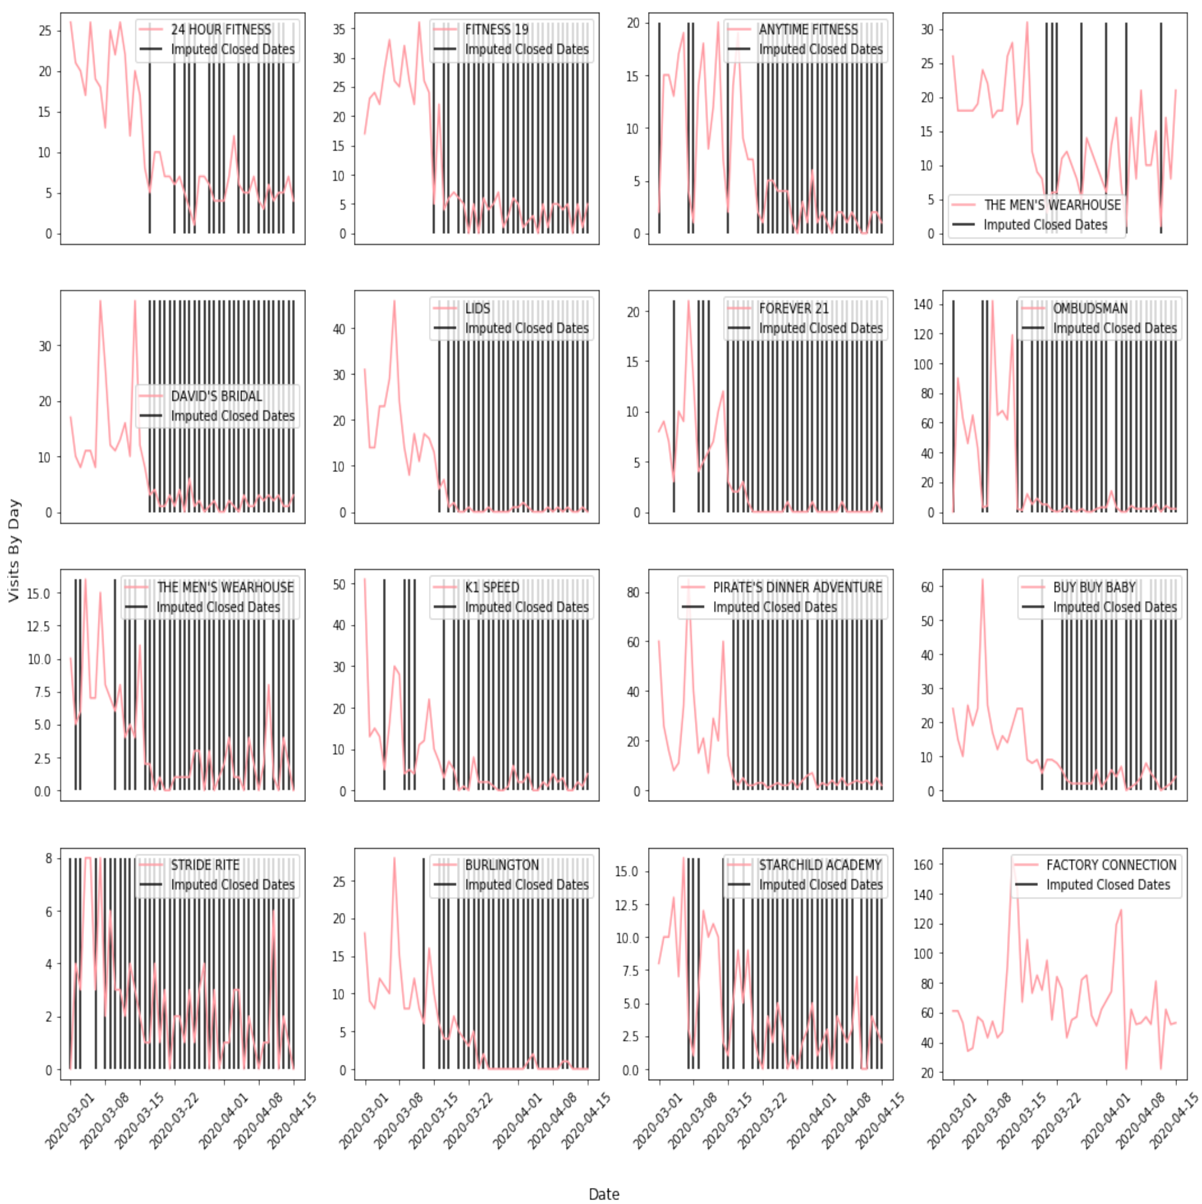

Finally, to further establish the robustness of our main findings, we re-estimated our main models by using two simple measures for whether an establishment is open or closed. Table A3 shows the results. In Model 13, we use the ratio of the number of visitors on a given date in March or April to the February average for the same day of the week. We use 20% as the threshold to label an establishment as open. In other words, we simply identify an establishment as open if $R_o > 0.2$. In Model 14, we count the number of visitors on a given day. Places with five or more visitors are identified as open. These models show that our results are robust to alternative ways of defining establishment closure.

Table A3: IV Regression of Community Est. Open using alternate methods to determine if an establishment is open or not

	Model 13		Model 14	
	First Stage	IV	First Stage	IV
Prop. Branch Est. Open$_{t-1}$		0.160^{***}		0.110^{***}
		(0.008)		(0.007)
National Chain Opening Exposure$_{t-1}$	0.988^{***}		0.981^{***}	
	(0.013)		(0.013)	
Avg. February Traffic	0.000	-0.001^{**}	0.000	0.002^{*}
	(0.000)	(0.000)	(0.000)	(0.001)
Prop. Devices At Home	0.038^{*}	-0.135^{***}	0.010	-0.179^{***}
	(0.019)	(0.017)	(0.016)	(0.014)
Fixed Effect Zip	Yes	Yes	Yes	Yes
Fixed Effect NAICS × Date	Yes	Yes	Yes	Yes
Fixed Effect County × Date	Yes	Yes	Yes	Yes
Observations	7,939,548	7,936,243	9,688,964	9,688,964
R2	0.600	0.255	0.629	0.314
Adj. R2	0.596	0.247	0.626	0.308
Number of Groups: Zip	11,879	11,879	11,879	11,879
Number of Groups: NAICS × Date	1,075	1,075	1,125	1,125
Number of Groups: County × Date	68,477	68,477	70,738	70,738

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: In Model 13 and Model 14 we use alternate definitions to determine whether an establishment is closed or open. In Model 13 establishments are considered *Open* if the number of customer visits is greater than 20% of the average number of customer visits in February on the same day of the week. In Model 14 establishments are considered *Open* if the number of customer visits is more than five.

Appendix B: Robustness Checks

As mentioned in the main text, there are a few threats to the validity of our instrumental variable. First, for our instrument to be valid, there should be no relationship between the instrument and local business and disease conditions. We report a series of analyses to get at this question in Figure B1, which appears in Appendix C. We created 20 equally spaced bins of our instrument and evaluated the relationship of the instrument with zip code characteristics such as mean income, income inequality, and racial composition as well as COVID-19 cases at the county level. As Figure B1 shows, there seems to be no systematic relationship between the instrument and these zip code and county level covariates, thus offering reassurance about the validity of the instrument.

In addition to local disease conditions, a second plausible example of common exposure includes similarity in customer types. The non-random location choices of chain and community establishments may expose them to similar customer types, and similarity in closing decisions informed by customer type could lead to the temporal clustering of closing behavior. For example, consider the case of Morton's Steakhouses that tend to be in wealthy neighborhoods with an older clientele. Community establishments that compete with Morton's are likely to have a similar customer base. If Morton's, recognizing the risk profile of their customers, chooses to close, nearby restaurants might also make the same calculus and close, leading to a false impression of social influence.[3] Even though this concern requires considerable similarity in customer type both within the chain and between the chain and the community establishment, we conduct a robustness check in which we limit our sample to community establishments that serve very different customers than do nearby chain establishments to evaluate this concern directly.

[3] We are grateful to an anonymous reviewer for pointing out this possibility and for pushing us to run analyses to help rule it out.

Figure B1: Plots of Census Covariates, COVID Cases and National Chain Opening Exposure (Instrumental Variable).

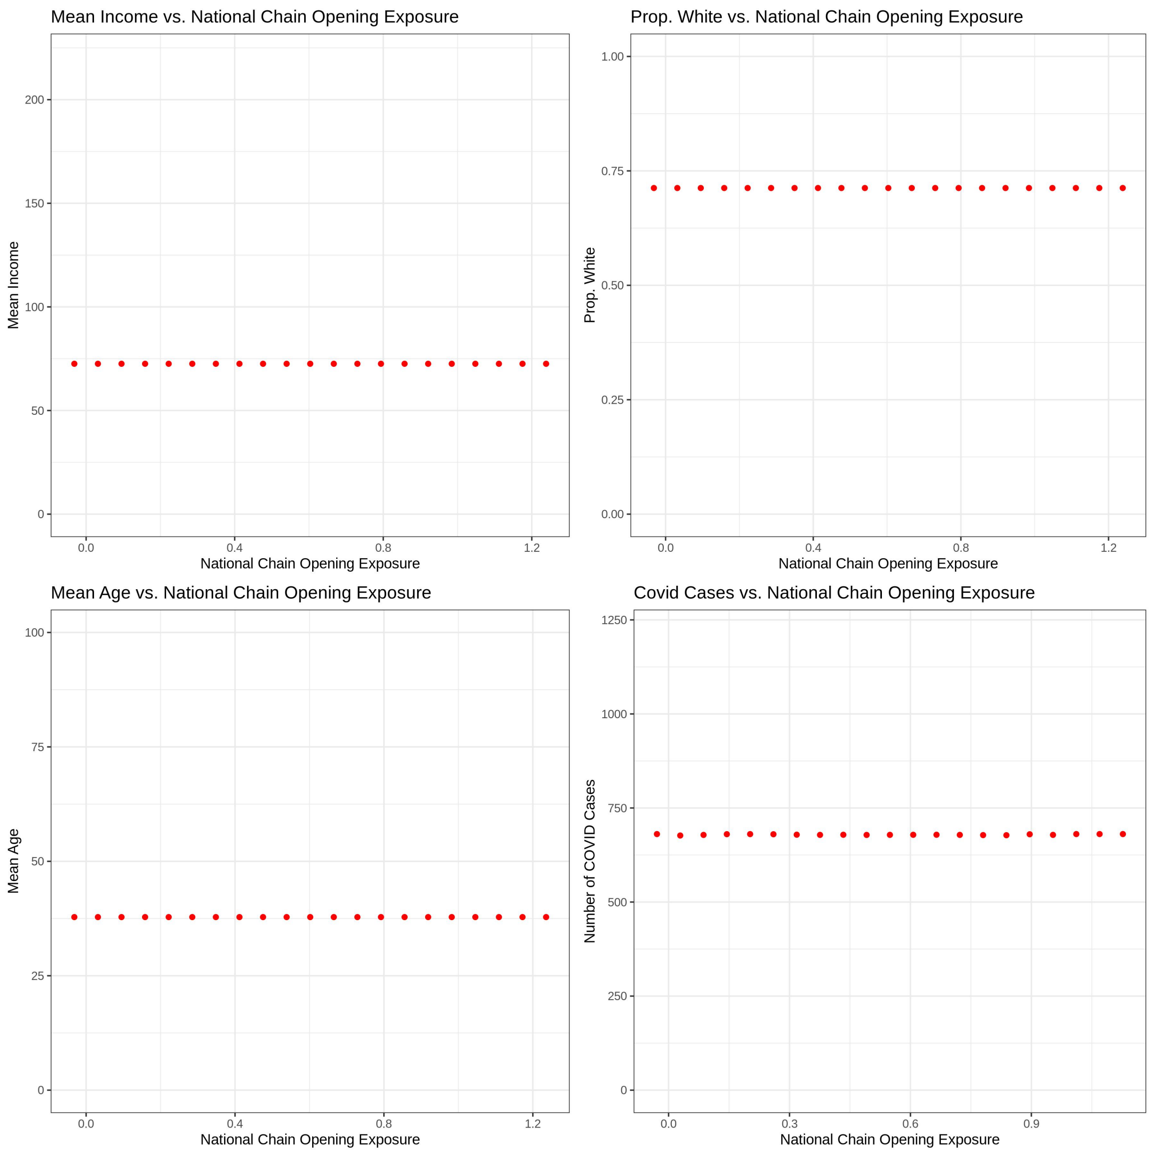

Note: This figure shows bin-scatter plots of Census covariates and COVID cases at the Zip level against the instrumental variable (*National Chain Opening Exposure*). *Mean Income* is the average income within a zip code. *Prop. White* is the proportion of white residents within a zip code. *Mean Age* is the average age of the residents in the zip code. *COVID Cases* is the number of reported COVID cases in a county on a day. The census covariates and COVID cases are uncorrelated with the instrument variable.

Table B1: Robustness Checks

	Model 7 (Excl. Clustered Est.)		Model 8 (High Compliance Brands)		Model 9 (Dissimilar Customers)		Model 10 (Excl. Border Est.)	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Prop. Branch Est. Open_{t-1}	0.026^{***}	0.054^{***}	0.065^{***}	0.187^{***}	0.031^{***}	0.079^{***}	0.035^{***}	0.092^{***}
	(0.002)	(0.008)	(0.004)	(0.015)	(0.003)	(0.011)	(0.002)	(0.009)
Avg. February Traffic	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Prop. Devices At Home	-0.302^{***}	-0.298^{***}	-0.187^{***}	-0.179^{***}	-0.170^{***}	-0.168^{***}	-0.160^{***}	-0.158^{***}
	(0.017)	(0.017)	(0.027)	(0.025)	(0.024)	(0.023)	(0.018)	(0.018)
Fixed Effect Zip	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effect NAICS×Date	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effect County×Date	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,225,808	5,225,808	6,986,468	6,986,468	3,805,101	3,805,101	6,539,134	6,539,134
R2	0.313	0.313	0.331	0.329	0.327	0.326	0.316	0.316
Adj. R2	0.302	0.302	0.323	0.322	0.313	0.312	0.310	0.309
No. of Groups: Zip	11,600	11,600	9,548	9,548	10,416	10,416	7,881	7,881
No. of Groups: NAICS×Date	1,125	1,125	540	540	855	855	1,125	1,125
No. of Groups: County×Date	70,198	70,198	65,761	65,761	65,363	65,363	52,333	52,333

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: This table restricts the sample using different conditions. Model 7 excludes all community establishments are located among a cluster of stores (Clusters were identified using DBSCAN. See appendix for details.) Model 8 restricts brand establishments to those brands that have highest compliance in opening/closing decisions across their establishment. Brand-level compliance is measured as the average daily variance in its establishments' opening status. Model 9 only includes local establishments that have low similarity to brand establishment in terms of origin of customers. In Model 10, we exclude all community establishments that are located in counties that are near the border between states.

The results are reported in model 9 of table B1. Establishments in our sample had visitors from different Census Block Groups (CBGs). We calculated visitor-similarity of two establishments as the cosine similarity between the origins of the visitors that visited these establishments. The sample used to estimate model 9 is then restricted to community establishments that are most dissimilar to the brand establishments in the same zip and NAICS code. We do this by computing the average brand visitor profile for every zip code and NAICS category, which is calculated as the average number of visitors from each CBG that visited any brand stores in the same zip code and NAICS category. We then compute visitor-similarity between all community establishments and the average brand visitor profile in the same zip code and NAICS category and exclude the community establishments with higher visitor-similarity (i.e., more than the median). This exclusion restricts the sample to community establishments that are dissimilar from their brand counter-parts. The results from this regression corroborate the findings reported in the main text and suggest similarity in customer base is unlikely to invalidate the instrument.

A third robustness check addresses concerns about the business model used by chains. Some brands rely on a franchise model, which reduces corporate control over the closure decisions of their chain establishments. Note that for our instrument to be valid, we do not need perfect compliance—just an increased likelihood of closure among chain establishments with national closure policies as compared to chain establishments without such a policy. As we have shown in the main text, there is a strong first-stage relationship between our instrument and the likelihood that any local brand establishment will close, helping to validate this assumption. However, to be sure that our results are not affected by chains with a large proportion of non-compliant franchisees, we implement an analysis in which we limit our sample of chain establishments to those that fall in the top quartile in terms of compliance with corporate guidelines. In other words, we only include those cases where the vast majority of establishments move in tandem in their closure decisions. Specifically, we calculate brand-level compliance as the average daily variance in its establishments' opening status. We include only the quartile of brands (70 out of 319) with the lowest variance levels. We assume that these brands are least likely to be franchise owned and we re-estimate our baseline models. The results of this regression are reported in model 8 of table B1 and are consistent with those reported in the main text.

Fourth, one might question our assumption that the instrument does not incorporate local business or disease conditions on the basis of geography. Recall that we only include those chains that have over 50 establishments distributed over at least 25 states and that we exclude the closing decisions of chain establishments in the focal state, making it unlikely that the closure patterns of chains incorporate local business and disease conditions. Some establishments, however, might be close to a state border and therefore be influenced by brand closing patterns (and infection rates) in the bordering state. Therefore, we conduct a robustness check excluding all counties that share a border with another state and estimate the baseline specifications. Specifically, we exclude all community establishments that are located in border counties. We mark a county as a border county if the centroid of the county is within 25 miles distance of a state border. We remove 1,118 counties and 75,092 number of community establishments. The results of a regression based on this reduced sample are found in model 10, table B1.

Fifth, economic mechanisms other than social learning might drive the hypothesized effects. For example, nearby chain and community establishments share support services (e.g., security workers), their employees share information, and, especially in malls, branded anchor stores drive traffic to smaller community establishments.[4] Such interdependencies between establishments may cause them to make similar decisions even in the absence of social learning. While we cannot rule all

[4]We thank an anonymous referee for raising this concern.

economic channels, many economic explanations rely on proximity between chain and community establishments. Therefore, we provide a robustness check in which we exclude chain and community establishments that are spatially clustered. Specifically, we do so by relying on the density-based spatial clustering of applications with noise (DBSCAN) algorithm that clusters establishments based on their geographic location. We tune the algorithm to detect groups of stores close to each other and detect 32,983 clusters in our data. Over 55% of all community establishments are in a cluster. We then remove these clusters (effectively, this removes malls and other areas of high spatial concentration of economic activity) and re-estimate our main regression model. The results are reported in model 7, table B1.

Across the four models, our estimates remains positive and statistically significant. In fact, in all models (except model 7), the OLS co-efficient is slightly larger than the baseline co-efficient, although the IV estimates reduce in size (except in Model 8). Combined, these sub sample analyses help build confidence in the validity of the baseline analysis.

As a final check to examine the concern that chain and community establishments are jointly exposed to the same local conditions and therefore exhibit similar behaviors, we also evaluate an event-study specification. In this analysis, the data are organized at the industry-zipcode-date level (rather than establishment-date level). Further, the key dependent variable is now *Post Brand Est. Closing*—a dummy variable that switches from one to zero after the first chain establishment within an industry has closed. We evaluate the likelihood that community establishments will close before and after the first chain establishment has done so. If business conditions or local infections are driving the closure of both chains and community establishments, we should see a rise in at least some community establishments closing *before* the first local chain closes. However, if our specification picks up the social influence channel, we should see that the effects arise only after the first chain establishment has closed and then increase in magnitude.

Figure B2: Event Study

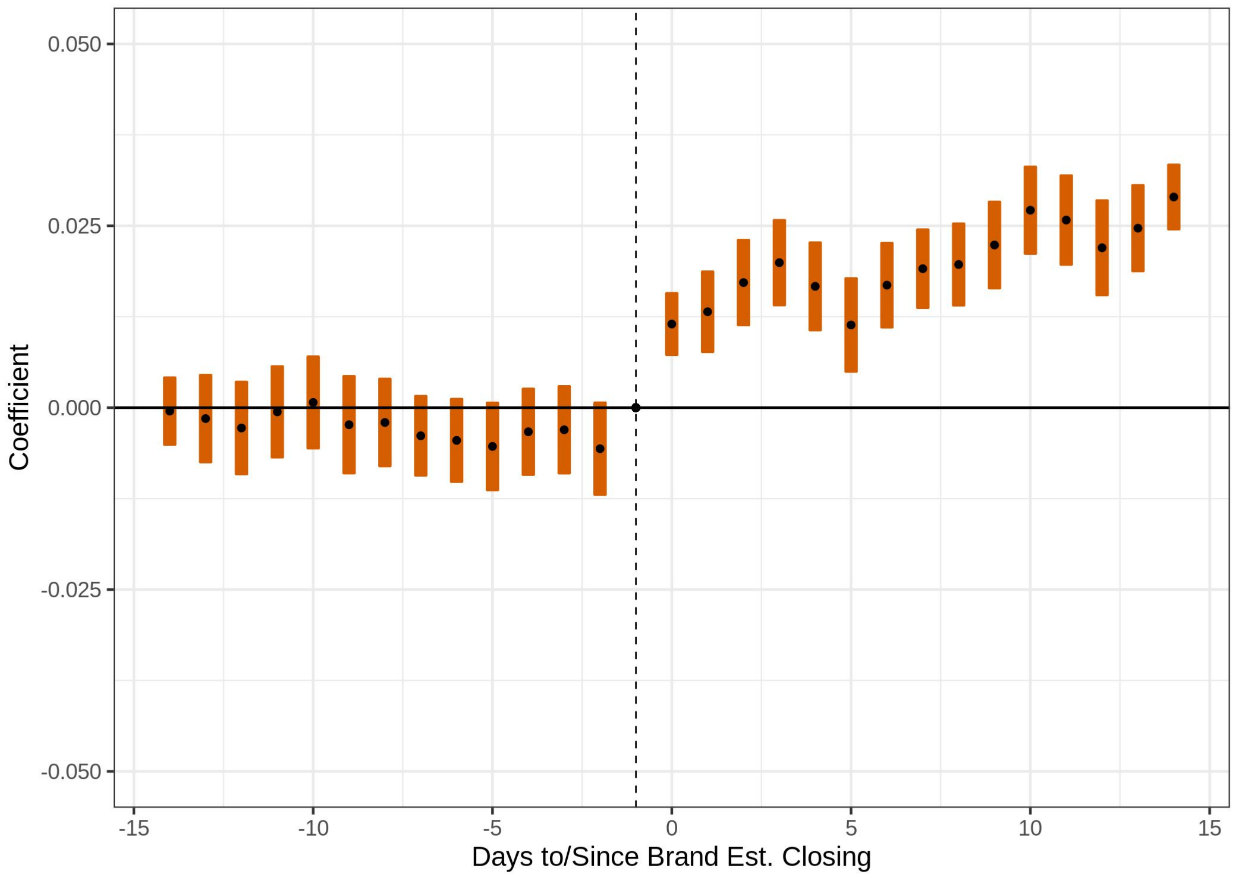

Note: This figure shows the event-study plot of the effect of closing of brand establishments on community establishments. We define establishments to be in a *closed state* if the establishment is closed for three consecutive weekdays. *Closed state* is an absorbing state - establishments once in a *closed state* continue to be in that state. For every Zip-NAICS pair, we find the first day that a brand store enters the *closed state* and use this date as the treatment date for that Zip-NAICS pair. The outcome is the proportion of community establishments closed in the same Zip-NAICS.

In Figure B2 we present coefficients from our event study specification at the industry-zipcode-date level. These are presented relative to the "focal date"—i.e., the date at which the first chain establishment in a given industry-zip closes. As is clear from this chart, the pre-trend is flat and overlaps with zero, suggesting that community establishments do not close ahead of chain establishments. This makes it unlikely that our main results are driven by strong common exposure effects, because one would likely see at least some pre-trend in that scenario. Rather, the graph shows that community establishments' closure follows soon after the closure of a nearby chain establishment. Further, this positive effect grows stronger and increases in magnitude about seven days after the first chain establishment shutters its doors. This event study chart provides an alternative research design to evaluate theoretical proposition and sheds further light on the dynamic unfolding of social influence. Appendix Table B2 provides more information and detailed estimates that underlie this chart.

Table B2: OLS Regression of Brand Est. Closing on Community Est. Closing

	Prop. Community Est. Closed			
	Model 17	Model 18	Model 19	Model 20
Post Branch Est. Closing	0.013^{***}	0.013^{***}		
	(0.001)	(0.001)		
Prop. Devices At Home		0.217^{***}		0.269^{***}
		(0.005)		(0.008)
Lag=-7			-0.003	-0.003
			(0.003)	(0.003)
Lag=-6			-0.004	-0.003
			(0.003)	(0.003)
Lag=-5			-0.004	-0.004
			(0.003)	(0.003)
Lag=-4			-0.002	-0.002
			(0.003)	(0.003)
Lag=-3			-0.003	-0.002
			(0.003)	(0.003)
Lag=-2			-0.006	-0.006
			(0.003)	(0.004)
Lag=0			0.012^{***}	0.013^{***}
			(0.002)	(0.002)
Lag=1			0.013^{***}	0.013^{***}
			(0.003)	(0.003)
Lag=2			0.018^{***}	0.018^{***}
			(0.003)	(0.003)
Lag=3			0.021^{***}	0.022^{***}
			(0.003)	(0.003)
Lag=4			0.018^{***}	0.018^{***}
			(0.003)	(0.003)
Lag=5			0.014^{***}	0.013^{***}
			(0.004)	(0.004)
Lag=6			0.019^{***}	0.018^{***}
			(0.003)	(0.003)
Lag=7			0.022^{***}	0.021^{***}
			(0.003)	(0.003)
Controls	No	Yes	No	Yes
Fixed Effect NAICS × Date	Yes	Yes	Yes	Yes
Fixed Effect County × Date	Yes	Yes	Yes	Yes
Observations	1,173,084	1,123,132	572,944	549,400
R2 (full model)	0.566	0.567	0.506	0.508
Adj. R2 (full model)	0.545	0.545	0.462	0.463
Number of Groups: NAICS × Date	825	825	818	813
Number of Groups: County × Date	52,107	51,874	45,904	45,713

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Note: This table presents the estimates of closing of brand establishment on community establishment using an event study design. The unit of analysis for this table is Zip-NAICS. For each Zip-NAICS pair, we define the treatment date as the first day that a brand establishment in the Zip-NAICS closes and remains closed for three consecutive weekdays. The outcome is the proportion of community establishment closed on a day.
Model 17 and 18 use 14-days pre and post time period. Model 19 and 20 present estimates for 7-days pre and post treatment time period.