Online Appendix:

Disruption and Displacement in Health Care: Evidence from Nursing Home Closures

Andrew Olenski

A Additional Details

A.1 Approach to Measuring Quality

Nursing home quality of care is an inherently difficult object to measure, and the deficiency citations studied in this paper are only one possible metric. Approaches common to industrial organization, such as the use of revealed preferences in consumer demand to infer quality, are broadly ill-suited to health care markets due to the presence of asymmetric information (Arrow 1963). In nursing home markets, patient preferences are particularly difficult to interpret due to the number of agents involved in the nursing home decision (family members, hospital discharge planners, etc.) as well as selective admissions policies unobservably restricting choice sets (Gandhi 2020). Note that although it is common in the literature to employ staffing ratios (such as the number of registered nurse hours per resident day) as a proxy for facility quality, I do not do so here for data limitations. Prior to the introduction of the Payroll Based Journal data in 2016, staffing ratios were measured only once per year and known to be systematically over-stated by facilities (Geng, Stevenson, and Grabowski 2019). Closing facilities, facing financial distress, may have faced differential incentives to over-state their staffing ratios. Accordingly, I instead focus on the more reliable quality of care violations as my primary measure of quality, which I supplement with a standard risk-adjusted survival measure, described in Appendix C.

A.2 Identifying Exits

A common issue in the literature on provider exits is identifying whether a specific facility that exits the data actually shut down, or merely changed the provider identifier due to a merger, acquisition, or new certification (Carroll 2019; Joynt et al. 2015). Previous approaches in the literature on hospital closures have conducted manual searches to identify 'true' exits. Unfortunately, this approach is less feasible in the nursing home setting, as (1) there are about three times as many nursing homes as hospitals, (2) changes in nursing home ownership/name are much more frequent making manual searches more challenging, and (3) exits occur at an order of magnitude greater rate.

To identify nursing home exits, I construct a candidate list of exits by linking the termination dates in the Provider of Service files with the last year a facility is observed in the LTCFocus panel, and by restricting to facilities whose final observed year is within one year of its termination date. For these candidate closures, I then apply the Uber H3 hexagonal spatial index to assign each facility to a narrow tile of approximately 0.1 square kilometers.[19] A closure is 'confirmed' if there is no new facility operating in the tile in the subsequent year. This procedure leaves me with a final sample of 1,104 nursing home exits occurring over the period 2001-2014.

Of course, this procedure may be imperfect. For instance, any transcription errors in the address will result in inaccurate geocoding, which may erroneously lead to a facility being labeled an exit when it did not, though spot-checks and congruence with state-level reports suggests that this concern is minimal. Nonetheless, to the extent that my procedure identifies false closures, the estimated mortality effects will be attenuated towards zero. Moreover, this novel approach to identifying provider exits can be extended to other settings where similar issues arise, such as the literature on hospital closures.

19. I find very similar results when I expand the tile size to 1 square kilometer. Further details available at https://eng.uber.com/h3/.

A.3 Minimum Dataset 2.0 and 3.0 Transition

My analysis spans the transition from the Minimum Dataset 2.0 (MDS 2.0) to the Minimum Dataset 3.0 (MDS 3.0) in October 2010. This transition involved a substantial revision of the MDS data, including changes to which variables are assessed, and largely improved the veracity of the data. Many variables were added to the MDS 3.0 which would be worthy of independent study. For instance, 'patient-focused' variables such as reports of mental health became a primary focus. However, because these variables did not exist in the MDS 2.0, I am unable to examine them for the vast majority of my sample, which spans nursing home exits from 2001-2014.

Fortunately, the only MDS variable that is crucial for my analysis is the facility identifier (the CMS Certification Number), which is consistently defined across the transition. This allows me to identify which residents receive assessments in which facilities in each quarter, allowing me to seamlessly track patients as they move across facilities over time. My main outcome variables — mortality and hospitalization — are measured using Medicare enrollment and claims data, which are consistently defined.

Note that my results do include other measures derived from the MDS. The patient covariates used in my primary results – the variables contained in X_i in equation (1) – includes only variables that are common to both the MDS 2.0 and MDS 3.0. Specifically, this includes the set of diagnosis codes that are common to both datasets, as well as indicators for whether the patient required help with any of the activities of daily living. To address the concern that coding practices for these variables may have changed over time, I engage in several robustness exercises. First, I examine the sensitivity of my results to excluding these variables entirely (Appendix Figure D.6). Second, I examine the sensitivity of my results to an alternative set of risk-adjustors that are derived from Medicare claims, and hence not subject to this transition (Appendix Figure D.7). Reassuringly, both exercises yield nearly identical results to my primary analysis.

A.4 Further Details on Matching Procedure

Matching is done using observed facility characteristics from the year prior to closure. Prospective (control) matches are evaluated by their distance from each treated facility. All matching variables come from the LTCFocus annual panel of nursing homes. Only freestanding (i.e., non-hospital based) nursing homes are considered.

Prospective matches are required to match exactly on the following set of discrete variables: the presence of a specialty Alzheimer's unit, facility ownership (including indicators for both for-profit status as well as chain-ownership), whether the facility is designed as 'concentrated' or 'competitive' based on its Herfindahl-Hirschman Index (HHI), and finally the county in which the facility is located is assigned one of 6 classifications ranging from large central metropolitan area to noncore, which come from the 2013 National Center for Health Statistics Urban-Rural Classification Scheme for Counties.

The Mahalanobis distance is calculated between each prospective match and treated facility using following continuous variables: the shares of the facilty's residents covered by Medicaid and Medicare, respectively, the number of registered nurse/licensed practical nurse/certified nursing aide hours per resident day, the occupancy rate, and an index for patient acuity. Up to 4 matches are considered for each treated facility. Facilities with fewer than 3 viable matches are excluded. Matches that are located in the same county as the treated facility are excluded. Matches can also serve as controls for multiple treated facilities. All facilities are identified using their 6-digit CMS certification number (CCN), i.e. variables AA6B in the MDS 2.0, A0100B in the MDS 3.0, and PROV1680 in the LTCFocus data.

Time trends in each of the matching variables for the selected facilities are shown in Appendix Figure D.2. The trends are smooth and show no evidence of sharp changes immediately preceding or following the reference facility closure. Note that these data are measured annually, so to ensure sufficient pre- and post-period data I restrict the sample to only exits occurring between 2004 and 2015.

B Cumulative Mortality Simulation

The primary object of interest in the paper is the effect on cumulative mortality (ΔM_t, described in equation (2)). Following Deryugina and Molitor (2020), I compute ΔM_t indirectly using estimates of the mortality hazards β^τ, estimated via equation (1). One natural question is: why not estimate the model using cumulative mortality directly? The answer is that direct estimation of the cumulative effect on mortality is infeasible in the presence of differences in baseline mortality rates, which necessitates a violation of the parallel trends assumption.

Intuitively, because cumulative mortality converges to one for all groups, any differences in baseline mortality *requires* that cumulative mortality between two groups must not move in parallel. Mortality hazards, however, need not converge to one, and so the parallel trends assumption can be satisfied when considering per-period rates.

To illustrate this point, I conduct a simple simulation exercise. In the simulation, there are two groups of individuals: treated and control. Those in the treated group experience a shock to their per-period mortality hazard, while individuals in the control group do not. This shock resembles the pattern estimated in the main text. In the baseline simulation, all individuals face the same baseline 7.06% mortality rate, the same as the treatment group as reported in Table 1. In the second simulation, the control group faces a slightly higher 7.42% mortality rate, as in Table 1. The aim is to explore how the two approaches — inferring the cumulative mortality effect indirectly, as I do in the paper vs. calculating the 'empirical' difference in cumulative mortality between treatment and control groups — perform under each of these scenarios. I simulate the mortality of 1 million individuals over 20 periods, with treatment occurring in period 6.

The results are shown in Figure B.1. The value of the shock (the change in per-period mortality hazards, analogous to β) is plotted in each panel. This pattern is chosen to mirror the results in the paper: the treatment group experiences a large short-run increase in mortality risk, followed by a long-run decline.[20] By construction, there is no violation of parallel trends.

The cumulative mortality effect ΔM_t implied by the values of β is also shown in each panel. Notably, these estimates are the same in both scenarios. Even in the presence of different baseline mortality rates, the method correctly shows no evidence of diverging pre-trends, and correctly calculates the implied effect on cumulative mortality for each period. This follows the standard assumptions of difference-in-differences, that 'balance-in-levels' is not necessary in the dependent variable, only parallel trends.

The 'empirical cumulative mortality' effect, which is the difference in survival curves for the two groups, is also calculated. In the first panel, where there is no difference in baseline mortality rates, the implied and empirical cumulative mortality effects are nearly identical. In contrast, in the second panel, the 'empirical' mortality effect shows immediate divergence between the treatment and control groups, even though there is no parallel trends violation coming from the treatment effect. Instead, this reflects the difference in baseline mortality rates between the treatment and control groups.

This inability of the empirical cumulative mortality effect to differentiate between unequal baseline rates and treatment effects precludes estimation of the effect on cumulative mortality directly, and gives rise to the method employed in the paper. Table 1 in the main text illustrates that there are indeed differences in baseline mortality rates between treatment and control groups, motivating the use of this approach rather than direct estimation of cumulative mortality effects.

20. That is, the hazard effects are briefly large and positive and then small and negative for the remainder of the panel.

(a) Identical Baseline Mortality Rates

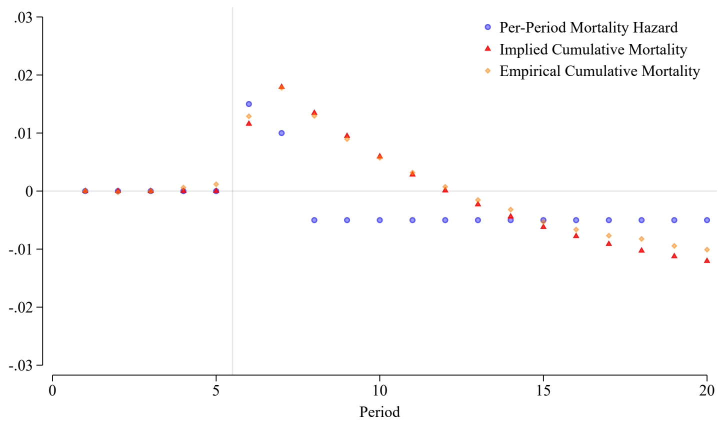

(b) Different Baseline Mortality Rates

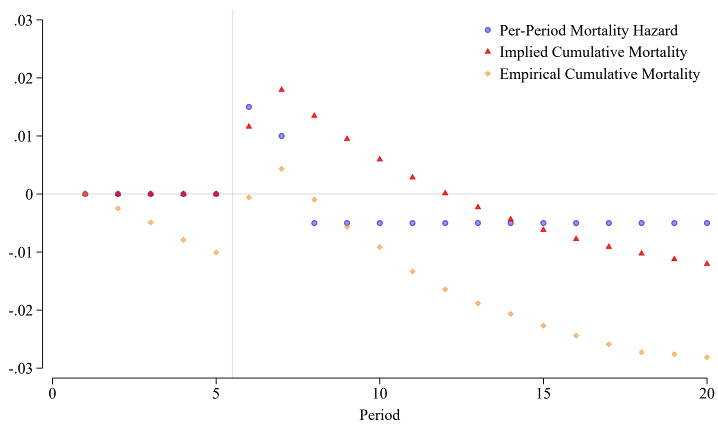

Figure B.1: Simulation Results

Notes: Figures present the results of a simulation exercise to explore the efficacy of inferring the effect on cumulative mortality, comparing treatment and control groups. The Implied Cumulative Mortality series reflects the cumulative mortality effect inferred using the method described in the paper. The Empirical Cumulative Mortality series calculates the effect taking the difference in survival curves between treatment and control groups. Panel (a) assumes that the treatment and control groups have identical baseline mortality rates. Panel (b) assumes that the treatment and control groups have different baseline mortality rates.

C Risk-Adjusted Survival Quality Measure

Models of risk-adjusted survival/mortality are common in the literature on health care quality estimation (e.g., Doyle Jr et al. 2015; Chandra et al. 2016). Following this literature, I construct a comparable measure of nursing home quality based on the risk-adjusted 90-day survival rates for all new patients (i.e., including both short-stay and eventual long-stay residents). While risk-adjusted survival from admission may have limited applicability to my sample of long-stay nursing home residents (who by definition have survived beyond 90 days after admission), it is a useful exercise to understand how the landscape of quality changes for short-stay residents after a facility closure.

To calculate risk-adjusted survival rates, I estimate a fixed effects regression of the form:

$$Y_i = \theta_{j(i),t(i)} + \alpha X_i + \eta_i$$

where i and j index admissions and facilities, respectively, t indexes a two-year bin (e.g., 2000 and 2001), Y_i is an indicator for whether the patient survives 90 days past admission and X_i is a vector of patient characteristics. The parameters of interest are the facility fixed effects $\theta_{j(i),t(i)}$, which can vary at the two-year bin-level. These terms capture the facility-specific component of the 90-day survival rate. Allowing the facility-specific effects to evolve over time is important, as one of the goals of this analysis is to measure how the distribution of quality in a market changes following a closure.

As in my main analyses, this sample is restricted to patients who are enrolled in the Medicare program, for whom I observe mortality in any setting (i.e., including at home or in the hospital). The vector of risk-adjustors X_i includes demographic information such as patient age, sex, and race as well as indicators for the presence of 26 comorbidities. My sample includes 19,070,127 nursing home admissions over the period 2000-2017. This sample excludes any admissions that occur within 90 days of a facility exit.

The results suggest that there is considerable dispersion in nursing home quality, consistent with existing estimates (Einav, Finkelstein, and Mahoney 2022; Cheng 2023). I find a one standard deviation increase in quality corresponds to a 5.38 percentage point increase in survival. The full distribution of resulting quality estimates is presented in Appendix Figure C.1.

Using this alternative measure, I examine the distribution of quality for several groups of facilities (Appendix Table C.1). Consistent with my primary results, exiting facilities appear to be low quality. Relative to other facilities in the market, exiting facilities have slightly lower mean quality (82.4% mean risk-adjusted survival, compared to 83.6% for the average facility in a market with an exit in that year-bin.) Notably, these means mask much of the differences in quality, which is in the far left-tail. The 5th percentile for closing facilities is only 70.6%, considerably lower than the corresponding percentile for other facilities in the market (74.5%).

Turning to the quality of 'absorbing' facilities to which displaced residents transfer, I again find results consistent with my primary analyses. Absorbing facilities appear to have much higher quality than exiting facilities, both at the mean and particularly in the left tail. The 5th percentile of quality in this group is 74.7%. Overall, absorbing facilities appear similar to the broader set of non-closing facilities, suggesting that patients are not reallocating to facilities that are particularly high- or low-quality. These results mirror my analysis of the distribution of quality-of-care deficiency citations: while there is much overlap in the distributions of quality, closing facilities are over-represented in the low-quality tail (Appendix Figure D.16).

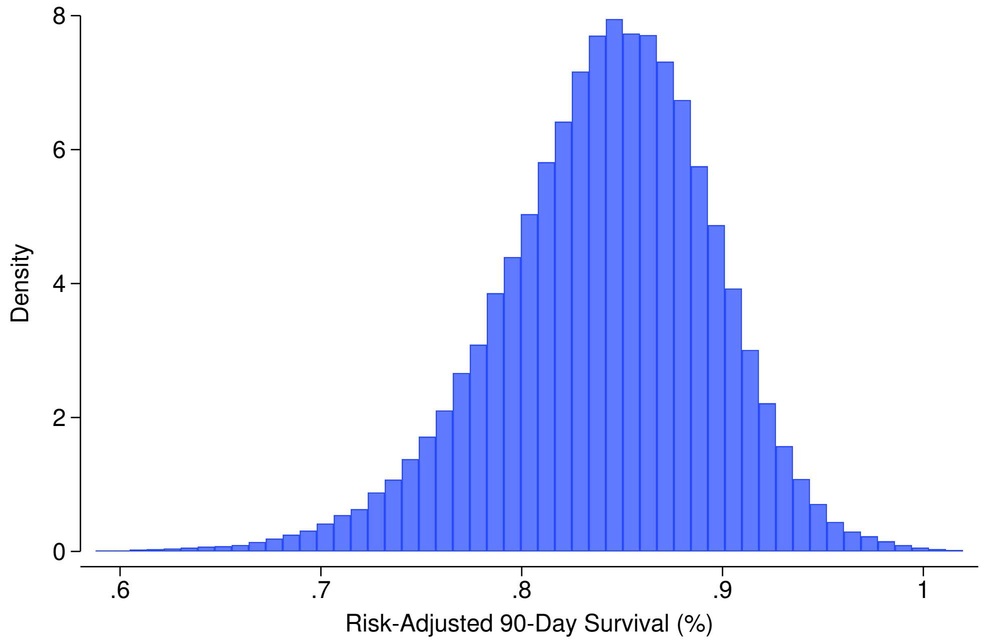

Figure C.1: Distribution of Risk-Adjusted 90-Day Survival Estimates

Notes: Figure presents the distribution of risk-adjusted 90-day survival estimates. Each observation is a facility-two year pair. Estimates are computed using all new nursing home admissions from 2000-2017 (excluding admissions within 90 days of a facility exit or the end of the sample). Estimates are trimmed at the 0.5th and 99.5th percentiles.

	Mean (1)	P5 (2)	P25 (3)	P50 (4)	P75 (5)	P95 (6)
Closing Facilities	82.42	70.63	78.46	82.76	87.12	92.24
Absorbing Facilities	83.53	74.69	80.35	83.82	86.98	91.27
Other Facilities in Market	83.57	74.45	79.93	83.88	87.22	92.42
Comparison Facilities	83.20	73.63	79.64	83.55	87.00	91.71

Table C.1: Distribution of Quality by Facility Status

Notes: Table presents the distributions of 90-day risk adjusted survival by facility status. Estimates correspond to the mean, 5th, 25th, 50th, 75th, and 95th percentiles. Risk-adjusted survival is computed using all nursing home admissions from 2000-2017.

D Additional Tables and Figures

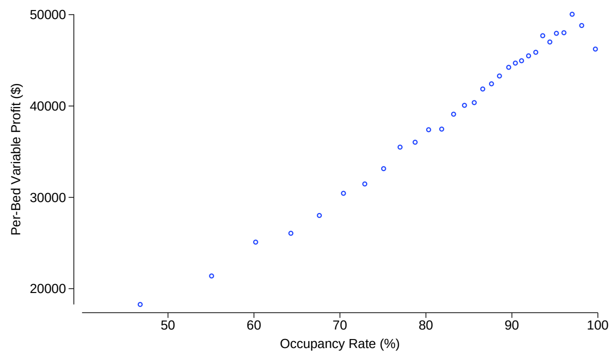

(a) Variable Profits by Occupancy Rate

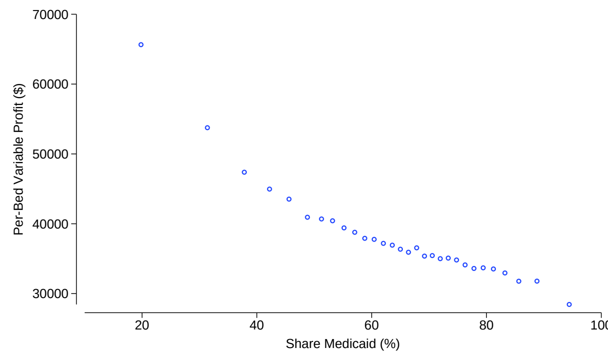

(b) Variable Profits by Medicaid Share

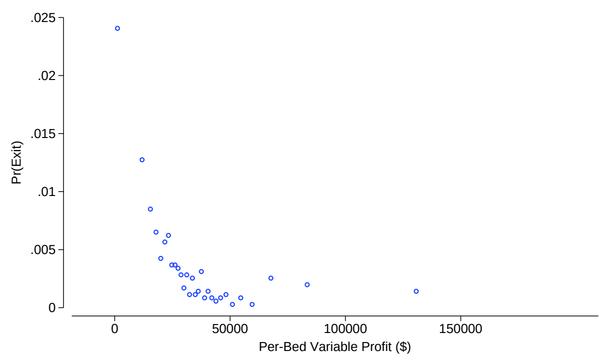

(c) Exit Rate by Variable Profit

Figure D.1: Determinants of Nursing Facility Exits

Notes: Figures present binned scatterplots of facility-year variable profits, exit probabilities, occupancy rates, and shares of patients whose stays are funded by Medicaid. Data on occupancy and Medicaid shares come from the LTCFocus.org database, while data on profits come from the Medicare Cost Reports, and span the period 2011-2019. Exits are defined in Section A.2.

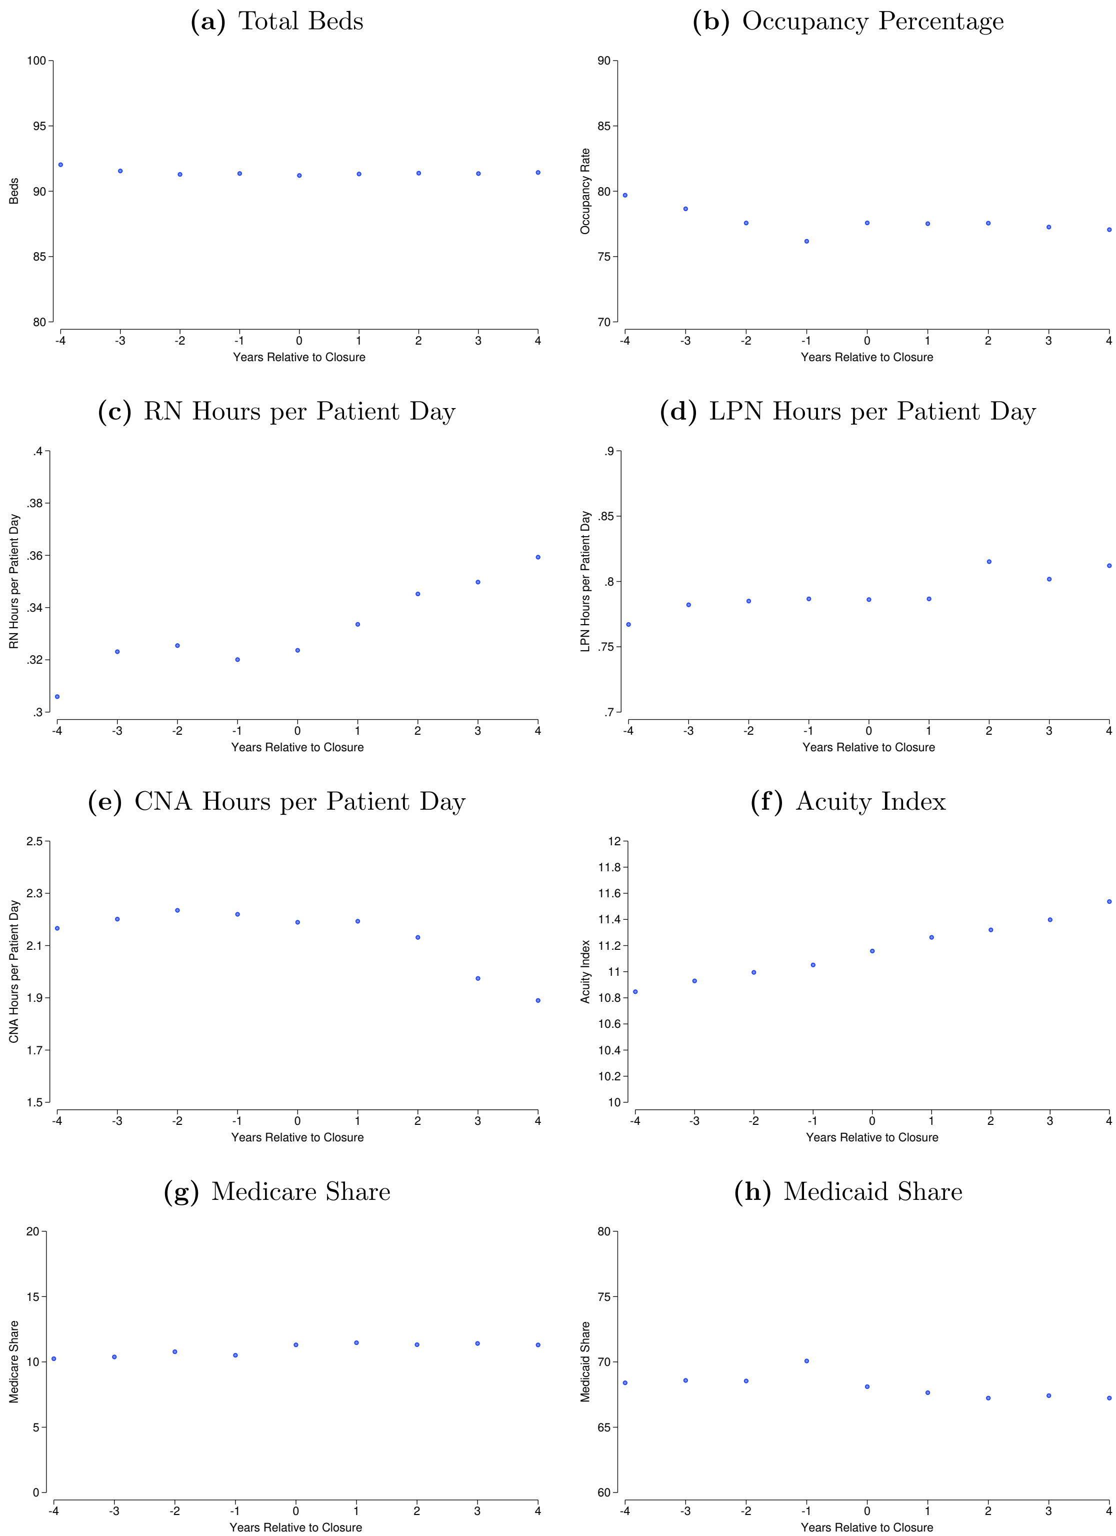

Figure D.2: Trends for Matching Variables

Notes: Figures present time trends in the continuous variables used in the matching procedure for the selected controls, relative to their reference facility's closure date. Data are drawn from the LTCFocus.org database.

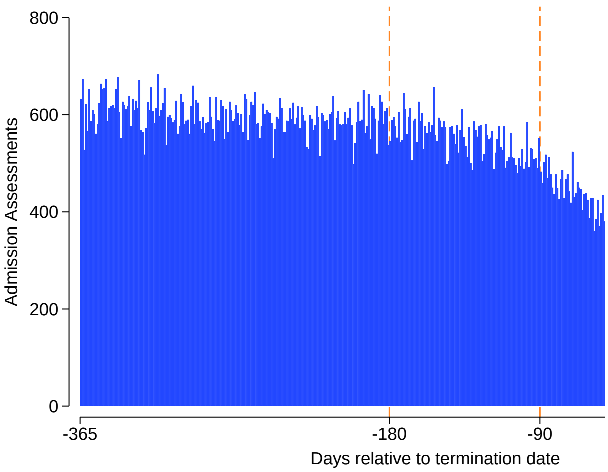

(a) Admission Assessments

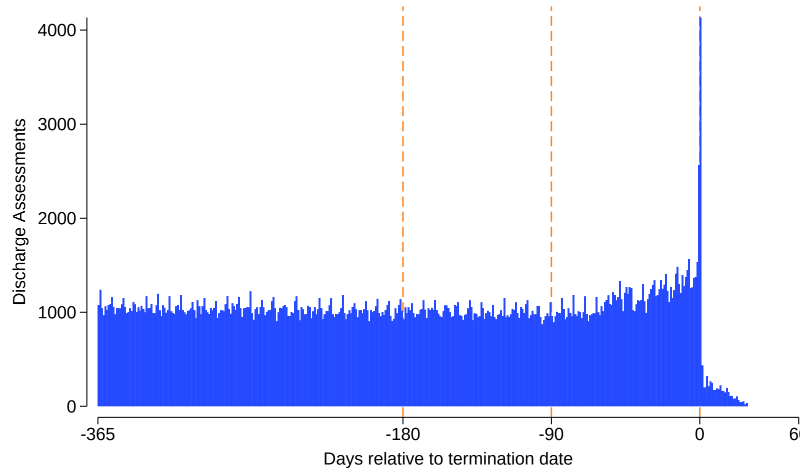

(b) Discharge Assessments

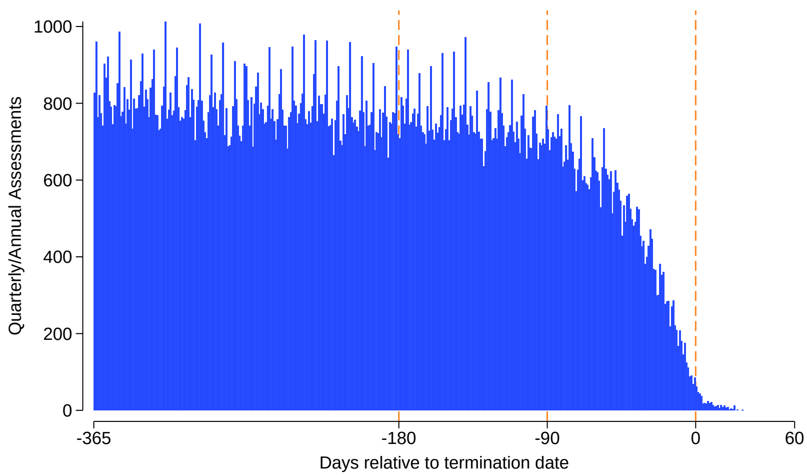

(c) Quarterly/Annual Assessments

Figure D.3: Counts of Assessments Relative to Exit Date

Notes: Figures present total daily counts of assessments across exiting facilities, by the date relative to its termination from Medicare and Medicaid. Figure (a) presents the counts of assessments corresponding to new admissions, which appears approximately stable until 90 days before the exit date, at which point they begin to taper. Figure (b) presents counts of discharge assessments, which also appear stable until 90 days before the exit date, at which point they rise sharply, with the largest spike occurring exactly on the date of exit. Figure (c) presents the counts of regular (quarterly or annual) assessments, which appear to follow a similar pattern as the admission assessments.

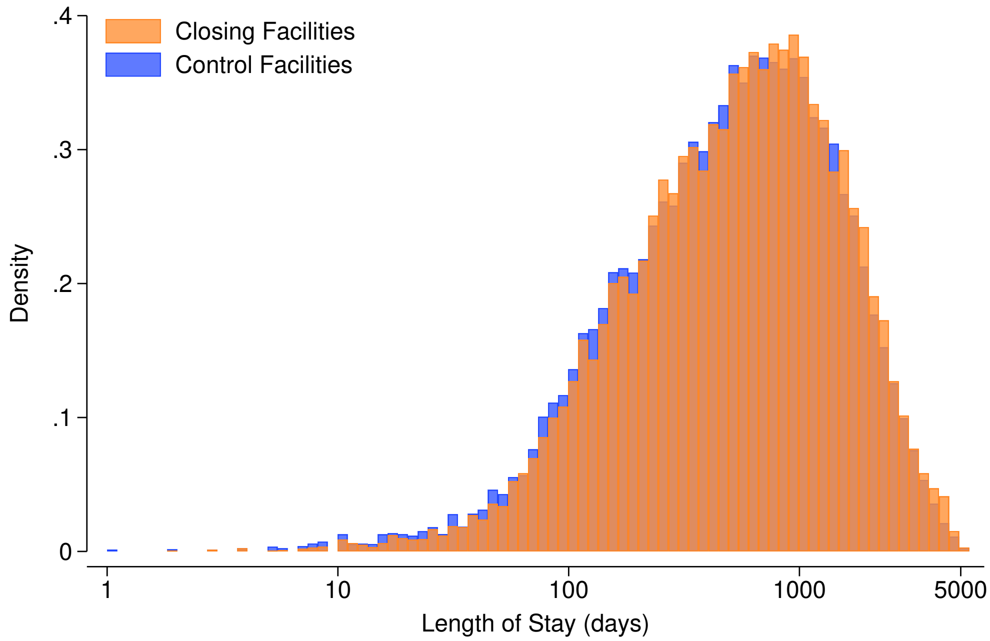

Figure D.4: Distribution of Length of Stay Prior to Closure

Notes: Figure presents the distribution of patient length of stay prior to facility closure. Histogram presented separately for patients in closing facilities and their matched controls.

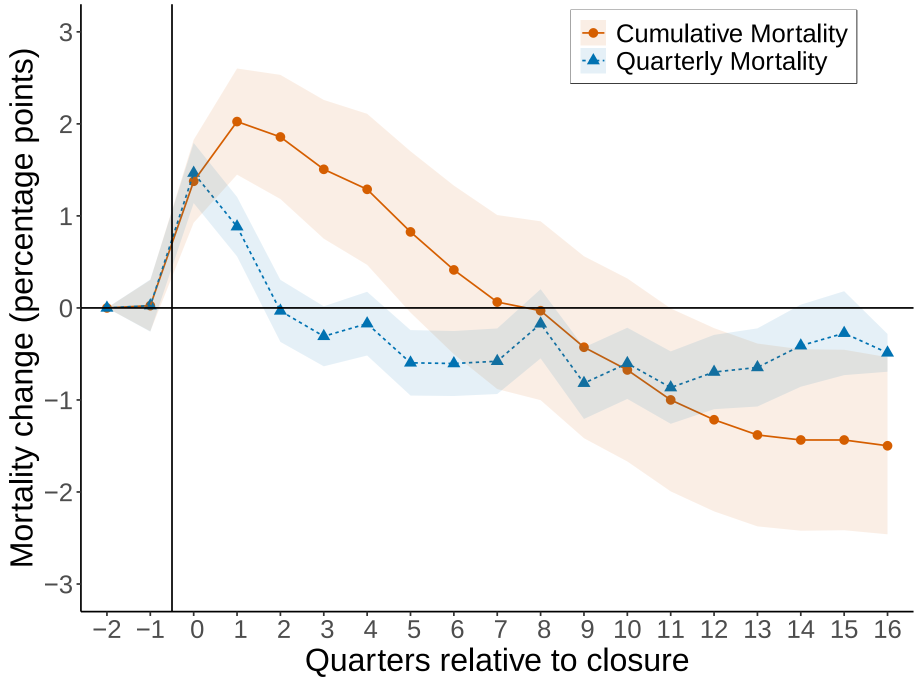

Figure D.5: Mortality Estimates with Extended Follow-Up

Notes: Figure presents estimates from a modified version of equation (1) along with the cumulative mortality estimates (ΔM_t) for the baseline resident cohort, that extends the follow-up period from the baseline 12-quarter to a longer 16-quarter window. The sample is restricted to facilities that closed prior to 2014 to ensure a balanced follow-up period.

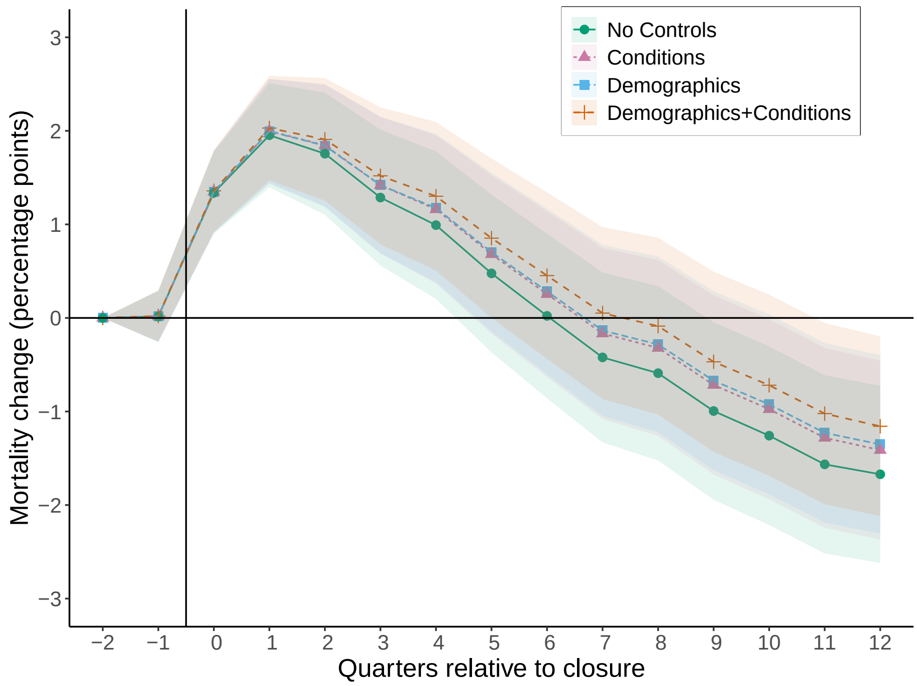

Figure D.6: Test of Coefficient Stability

Notes: Figure presents several estimates of the cumulative mortality effect (ΔM_t) for the baseline resident cohort, allowing for differing levels of controls X_i.

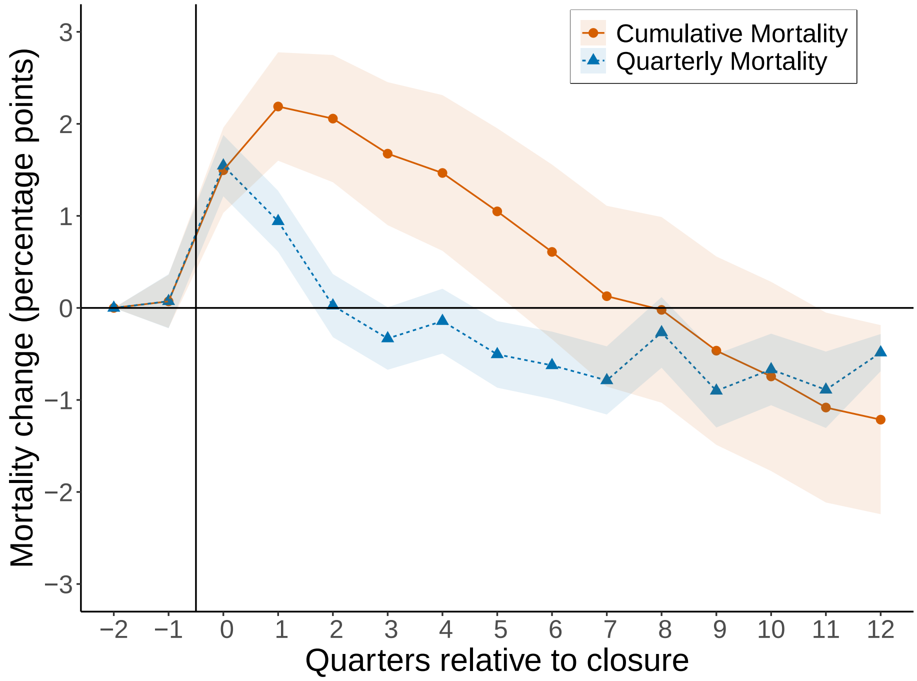

Figure D.7: Mortality Effect with Claims-Derived Controls

Notes: Figure presents estimates from equation (1) along with the cumulative mortality estimates (ΔM_t) for the baseline resident cohort, which in addition to the usual demographic variables, also includes a vector of 24 chronic condition indicators present in the Beneficiary Summary File. Because these codes are not well defined for Medicare Advantage patients, I restrict the sample to only fee-for-service Medicare enrollees.

(a) 2-Quarter Cohort

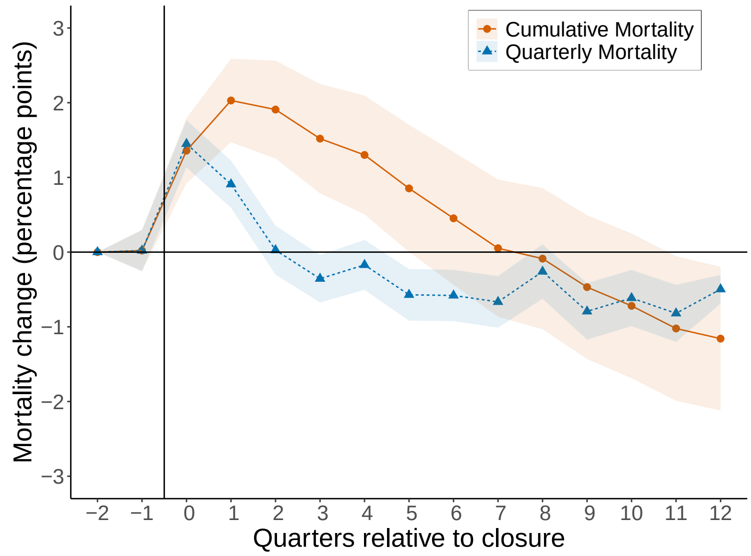

(b) 4-Quarter Cohort

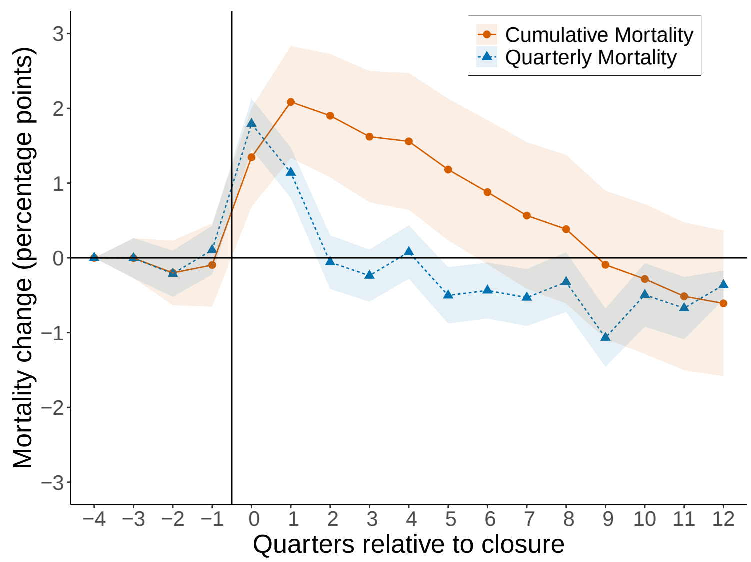

Figure D.8: Estimates Excluding Involuntary Closures

Notes: Figure presents estimates from equation (1) along with the cumulative mortality estimates (ΔM_t) excluding the 6.9% of facility closures that are associated with an involuntary termination from the Medicare/Medicaid programs.

(a) 2-Quarter Cohort

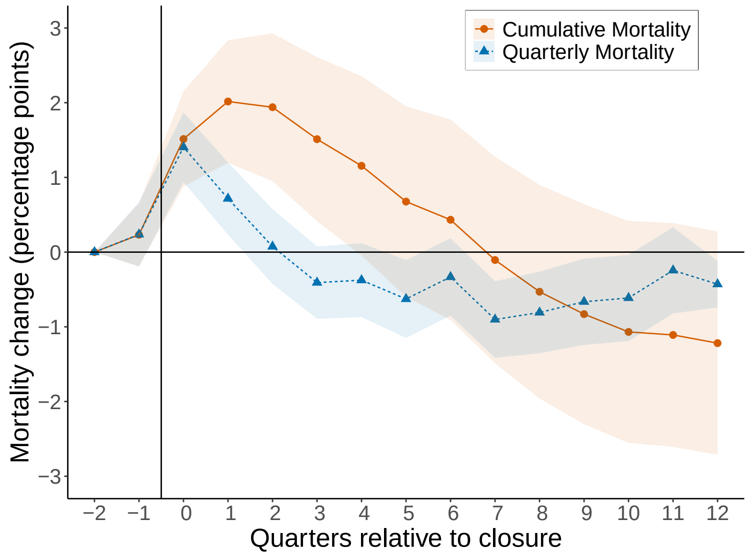

(b) 4-Quarter Cohort

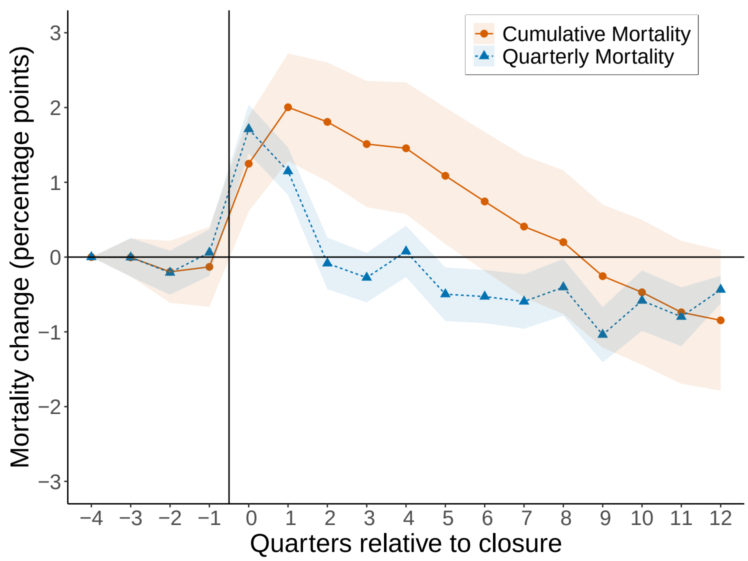

Figure D.9: Estimates from Alternative Matching

Notes: Figure presents estimates from equation (1) along with the cumulative mortality estimates (ΔM_t) for an alternative cohort of residents that come from incorporating quality-of-care deficiencies in the matching procedure. Sample is limited to 2007-2014 due to data restrictions.

(a) 2-Quarter Cohort: Competitive Counties

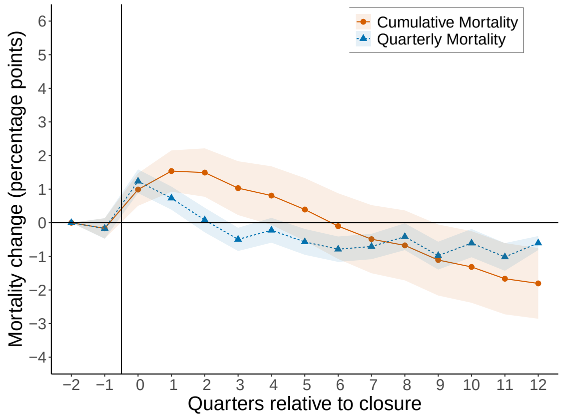

(b) 4-Quarter Cohort: Competitive Counties

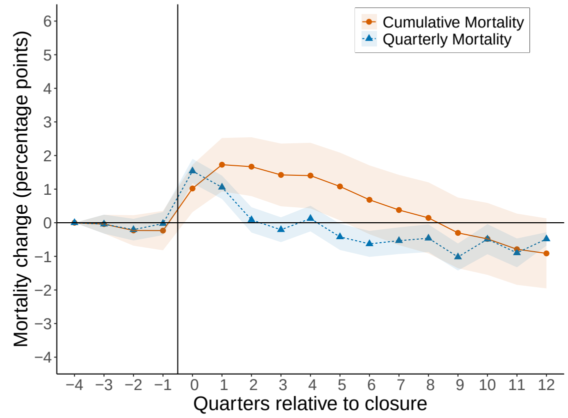

(c) 2-Quarter Cohort: Non-Competitive Counties

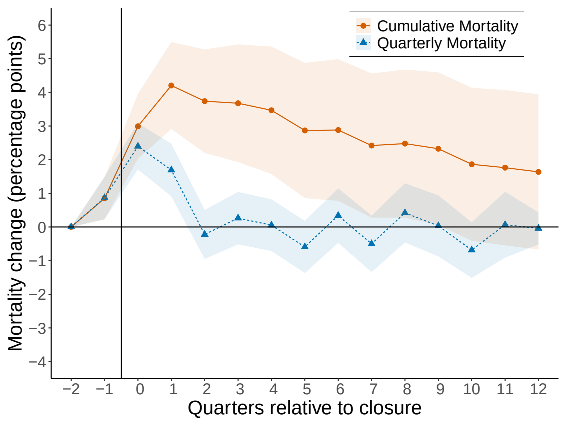

(d) 4-Quarter Cohort: Non-Competitive Counties

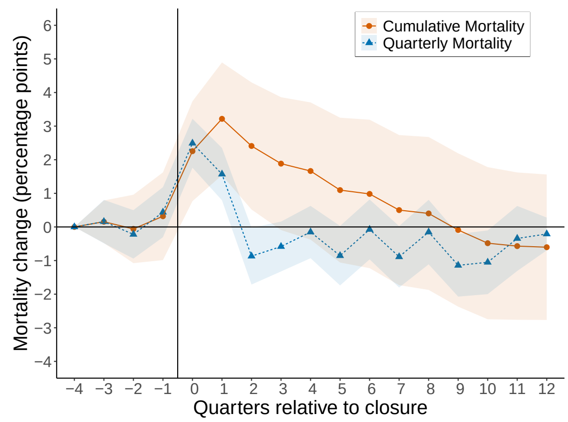

Figure D.10: Mortality Change by County-Level Market Concentration

Notes: Figures present estimates from equation (1) along with the cumulative mortality estimates (ΔM_t) for the baseline resident cohort by the level of pre-closure competition, using a county-level HHI measure. Given the larger market definition, I set the threshold for concentrated markets to those with HHIs above 2,500, which produces approximately the same share of facilities defined as concentrated as in the main definition.

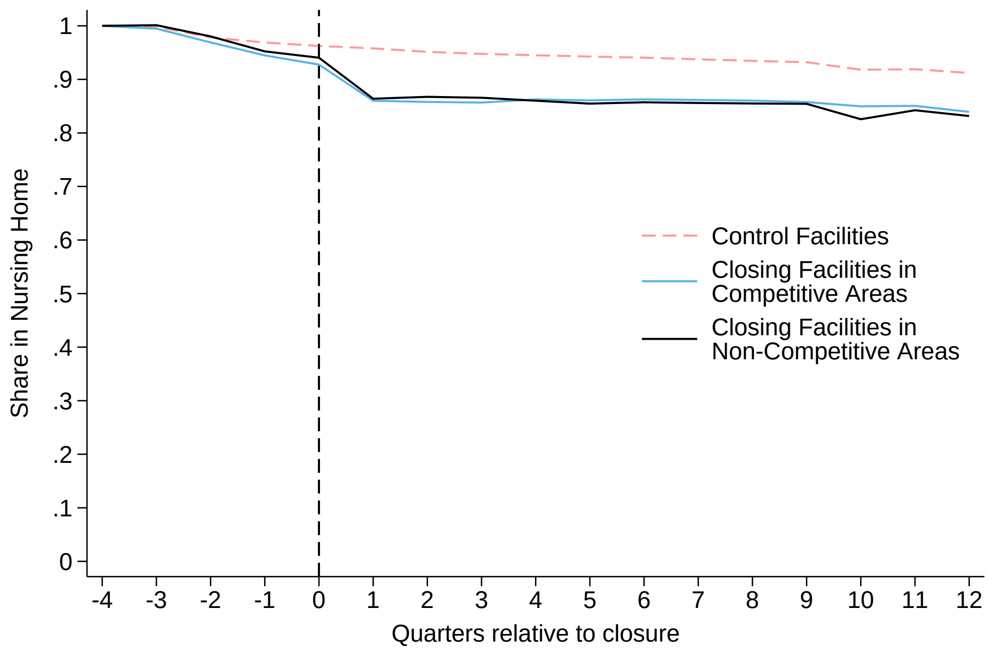

Figure D.11: Share of Surviving Cohort Still Present in a Nursing Home: 4-Quarter Cohort

Notes: Figures present the empirical share of residents who are still in a nursing home for the 4-quarter lookback cohort.

(a) 2-Quarter Cohort

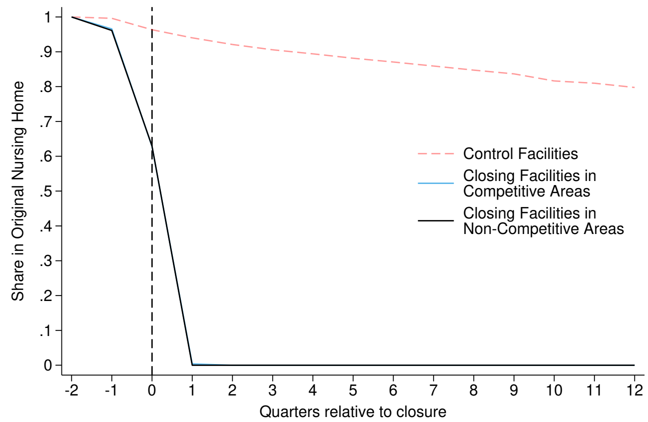

(b) 4-Quarter Cohort

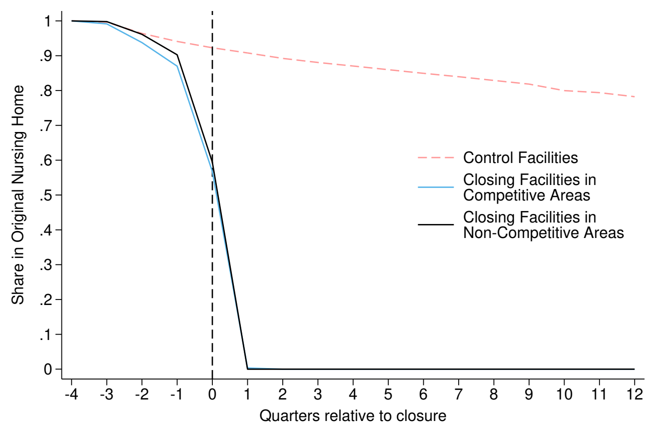

Figure D.12: Share of Surviving Cohort Present in Original Nursing Home

Notes: Figures present the empirical share of residents who are still in their original nursing home. Top panel presents estimates for the 2-quarter lookback cohort. Bottom panel presents estimates for the 4-quarter lookback cohort.

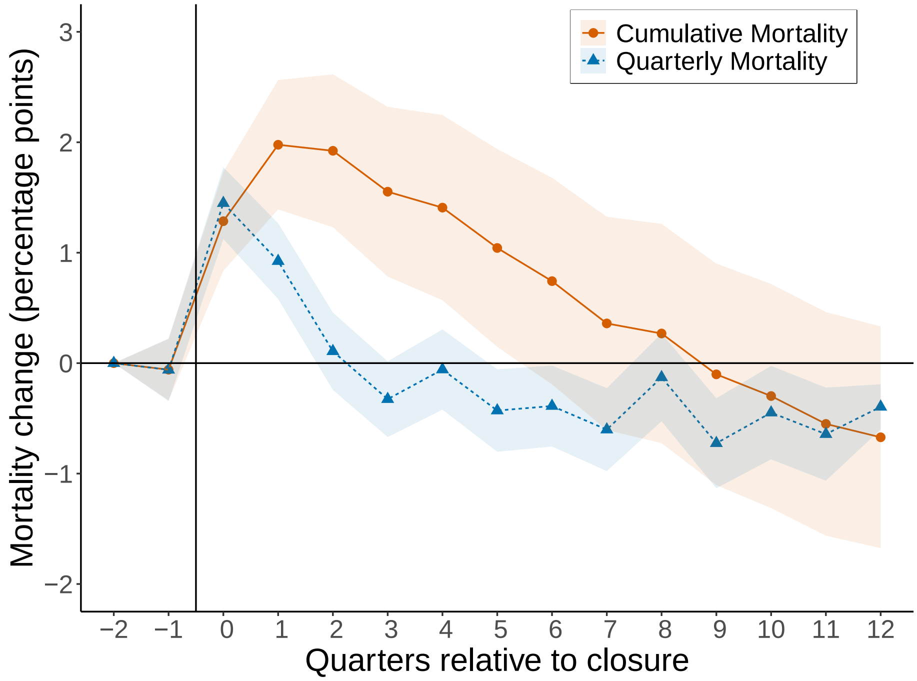

Figure D.13: Mortality Rate Relative to Closure

Notes: Figure presents estimates from equation (1) along with the cumulative mortality estimates (ΔM_t) for the baseline resident cohort, restricted to patients who are continuously present in any nursing home.

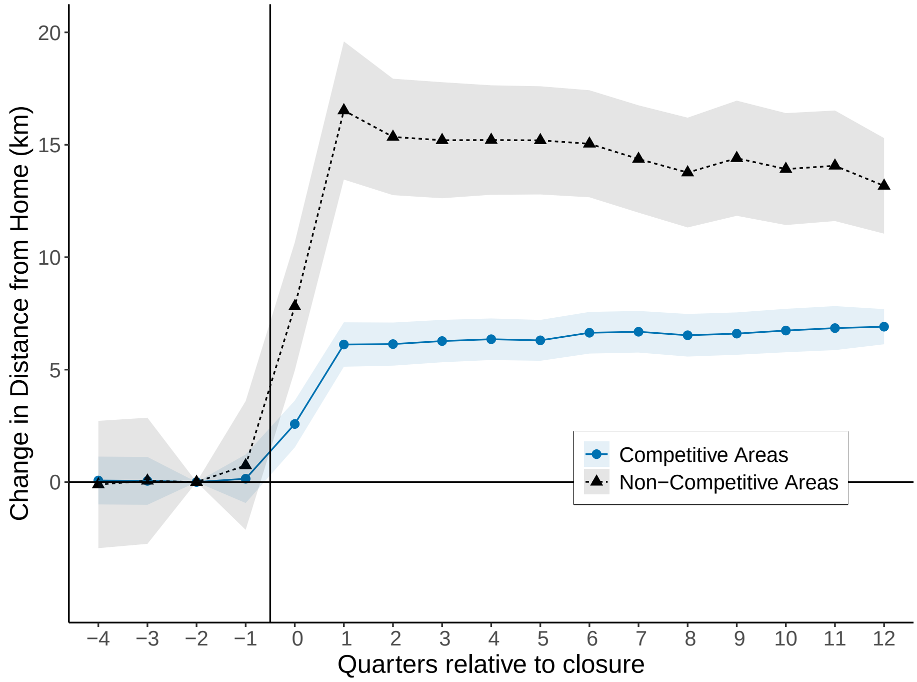

Figure D.14: Distance from home zip code

Notes: Figure shows how far patients are displaced following their nursing home closure, presenting β^{τ} estimates of equation (1) with the distance from the resident's home zip code to their current nursing home in each quarter as the dependent variable. Distance is determined using the resident's last 5-digit zip code (from the Medicare enrollment records) prior to their initial nursing home stay. Heterogeneous effects are estimated jointly, interacting the concentration measure with the relative time indicators. Patients who do not transfer to a new nursing home are excluded.

(a) Total Deficiency Citations

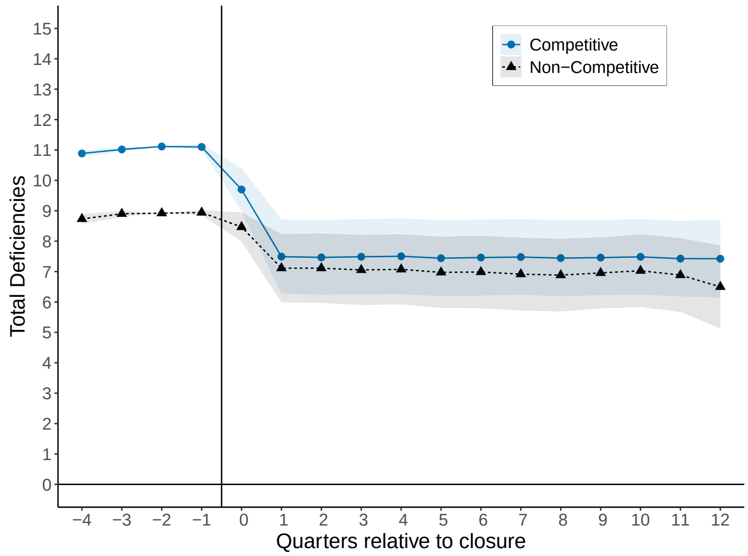

(b) Deficiencies for Actual Harm/Jeopardy

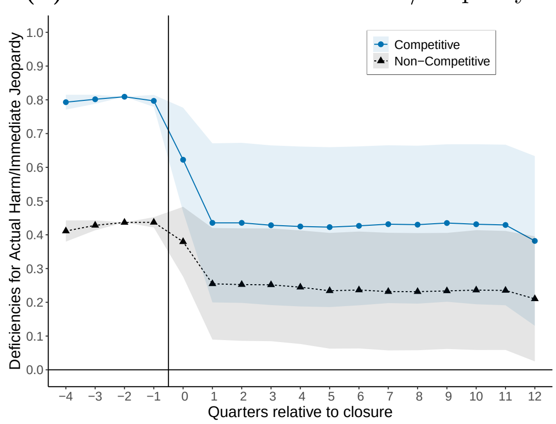

Figure D.15: Alternative Deficiency Citation Measures

Notes: Figure presents how patients reallocate following their displacement from a closing nursing home. The top panel reports the total deficiency citations of the resident's current facility. The bottom panel presents results restricting to only deficiency citations for actual harm or immediate jeopardy, the most severe categories. These figures are analogues to Figure 7c in the main text, which considers only quality of care violations. Patients who do not transfer to a new nursing home are excluded.

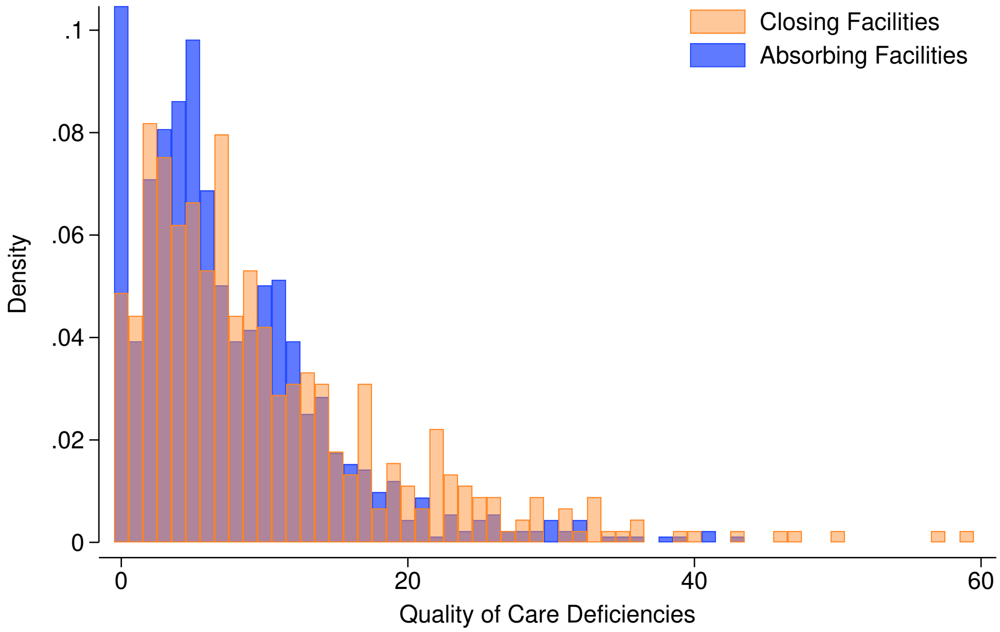

Figure D.16: Distribution of Quality of Care Deficiencies

Notes: Figure presents the distribution of quality-of-care deficiencies for each closing facility, as well as the set of 'absorbing' facilities, defined as the first facility a treated patient moves to following closure.

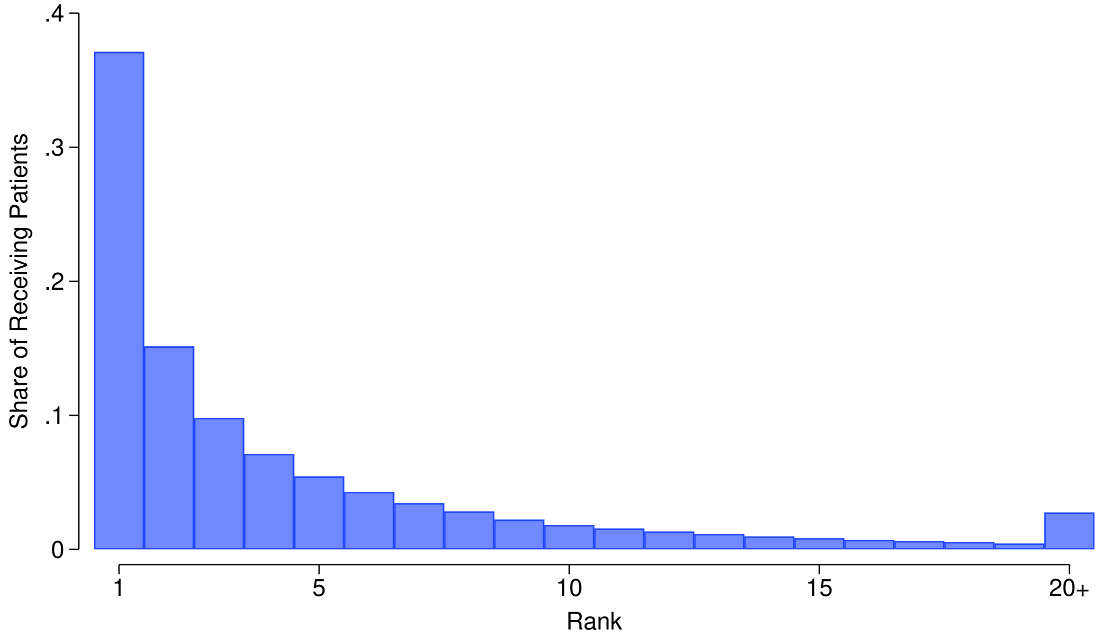

Figure D.17: Absorbing Facilities

Notes: Figure presents the distribution of absorbing facilities for residents displaced by nursing home closures. The figure shows the share of residents from the closing facility that move to the facility that receives the kth most residents. The figure shows that 37.1% of residents move to the most common receiving facility, and 52.2% move to one of the top two facilities. Figure is top-coded at the 20th ranked facility.

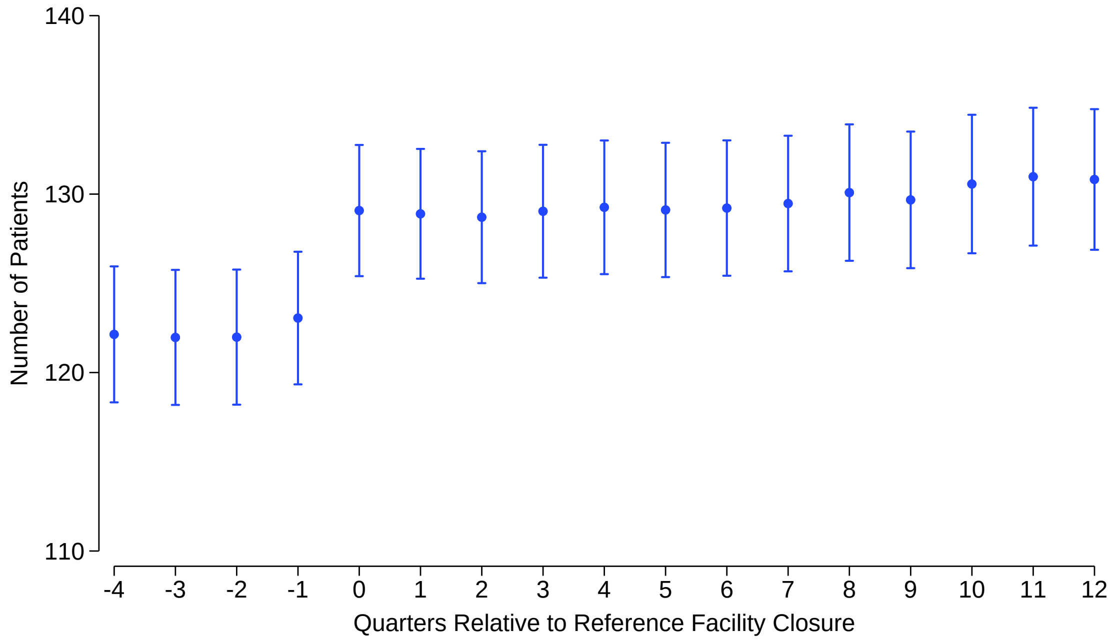

Figure D.18: Number of Patients in Absorbing Facilities

Notes: Figure presents the number of unique quarterly patients assessed in the absorbing facilities by quarter relative to the reference facility's closure date. Figure is restricted to the facilities that receive either the highest or second highest number of displaced patients.

Category	Short-Stay (1)	Long-Stay (2)
Total	13,280	15,641
Hospital Inpatient	4,140	4,852
Hospital Outpatient	1,846	1,770
Physician Office	1,736	1,661
Prescription Drug	1,432	1,410
Home Health	861	1,485
Other Inpatient	667	1,075
Hospice	183	305
Other	2,414	3,084

Table D.1: Average Annual Medicare Spending for Nursing Home Residents Prior to Admission

Notes: Table presents average annual Medicare spending by short-stay and long-stay patient status. Estimates calculated using the year prior to nursing home admission to avoid any differences in costs associated with nursing home use.

Data Source	Years	Relevant Data
MDS	2000-2017	Demographics, health measures, nursing home stays
BSF	2000-2017	Mortality dates, home zip codes, health measures, Medicaid eligibility
MedPAR	2000-2017	Hospital admission, Medicare coverage
LTCFocus	2000-2017	Nursing home characteristics
Deficiency Surveys	2006-2017	Deficiency citations
Medicare Cost Reports	2011-2017	Nursing home financials

Table D.2: Data Sources

Notes: Table lists each data source, the years spanned and the relevant data contained therein.

	Closed Firms (1)	Matched Firms (2)	All Non-Closed Firms (3)
Facility Characteristics			
Total Beds	84.6 (57.5)	94.6 (48.3)	110.5 (61.5)
Alzheimer's Unit, %	9.1 (28.7)	9.4 (29.2)	18.9 (39.1)
For-Profit, %	74.4 (43.7)	72.8 (44.5)	71.9 (45.0)
Concentrated, %	31.0 (46.3)	29.6 (45.7)	30.6 (46.1)
County Population, %			
Large Central Metro	24.5 (43.0)	25.6 (43.6)	21.3 (40.9)
Suburban	14.8 (35.5)	15.0 (35.7)	19.9 (39.9)
Small/Medium Metro	28.6 (45.2)	29.2 (45.5)	30.3 (46.0)
Rural	32.2 (46.7)	30.2 (45.9)	28.4 (45.1)
Patient Characteristics, %			
Occupancy	70.4 (19.1)	76.9 (16.1)	84.8 (14.3)
Medicaid	74.0 (20.1)	71.0 (18.7)	63.0 (22.0)
Private-Pay	18.2 (16.8)	19.0 (14.9)	24.7 (19.0)
Profit Margin, %	-8.8 (12.4)	-0.7 (9.9)	1.8 (9.2)
N	1,104	3,812	197,802

Table D.3: Facility Summary Statistics

Notes: Table presents summary statistics on the exiting facilities, their matched controls, and the universe of non-exiting facilities collected from LTCFocus.org and the Medicare cost reports. Standard deviations are reported in parentheses. Observations in columns (1) and (2) are drawn from the year prior to closure. Column (3) includes all observations for each non-closing facility. Because the distribution of exit years is not uniform, the observations in (3) are weighted to reflect the distribution of exit years, in order to facilitate comparison.