NORTHWESTERN UNIVERSITY

ESSAYS ON THE ECONOMICS OF HEALTH AND HOUSING POLICY

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of ECONOMICS

By

NICOLE FELLIN OZMINKOWSKI HOLZ

EVANSTON, ILLINOIS

JUNE 2023

ABSTRACT

ESSAYS ON THE ECONOMICS OF HEALTH AND HOUSING POLICY

NICOLE FELLIN OZMINKOWSKI HOLZ

This dissertation explores the economics of health and housing policies. In the first chapter, I discuss information disclosure policies in healthcare. Information disclosure programs can help consumers make better choices, but the consumers who respond the most to the information may not benefit the most or generate the most savings for firms designing the programs. I examine a disclosure program used by a private health insurer that highlights status of physicians based on two dimensions: quality and cost. Using a regression discontinuity design, I find that the higher physician status leads to 38% more new patients for physicians with that status (when compared to physicians with the lower status), and that the disclosure program's influence is stronger for younger patients. These young patients, however, may not benefit in terms of spending as much as other patients from being matched to higher-quality and lower physicians. Using a two-way fixed effects research design studying patients who switch physicians following physician exit (from retiring or moving), I find that switching to higher-quality and lower cost physicians leads to larger declines in spending for middle-aged patients than for younger

adult patients, with no evidence of adverse effects in terms of preventable emergency room visits. Collectively, these results indicate that targeting disclosure programs to middle-aged patients can achieve greater cost savings than web-based ratings systems that are disproportionately used by younger patients.

In the second chapter, I explore a question related to both health and housing: how rent control impact intimate partner violence (IPV). Policy advocates claim that one benefit of rent control may be decreased intimate partner violence. However, the theoretical effects of rent control on IPV are ambiguous. Rent control may lessen financial stressors within a relationship and decrease strain that leads to violence. However, it may make leaving the relationship more costly, shifting the bargaining power in the relationship and leading to more violence. My coauthor and I leverage the 1994 expansion of rent control in San Francisco as a natural experiment to study this question. This expansion created variation across zip codes in the number of rental units that were newly rent controlled. We exploit this variation in a continuous difference-in-difference design. We estimate an elasticity of -0.08 between the number of newly rent controlled units and assaults on women resulting in hospitalization. This effect translates to a nearly 10% decrease in assaults on women for the average zip code. This relationship is not explained by changes in neighborhood composition or overall crime, consistent with the effects being driven by individual level changes in IPV.

In the final chapter, I explore further consequences of rent control policies. Rent control policies seek to ensure affordable and stable housing for current tenants; however, they also increase the incentive for landlords to evict tenants since rents re-set when tenants leave. My coauthor and I exploit variation across zip codes in policy exposure to the 1994 rent control referendum in San Francisco to study the effects of rent control on eviction behavior. We find that for every 1,000 newly rent controlled units in a zip code, there were 12.05 additional eviction notices filed in that zip code and an addi- tional 4.6 wrongful eviction claims. These effects were concentrated in low income zip codes.

Acknowledgements

I would like to express my deepest appreciation to my chair, Matthew Notowidigdo, and to my committee: David Dranove, Dean Karlan, and Molly Schnell.

I am grateful for the help of countless others who provided input into these papers, including Desmond Ang, Lori Beaman, Kitt Carpenter, Marcus Dillender, Monica Garcia Perez, David Molitor, Scott Nelson, Ron Ozminkowski, Amanda Starc, and Daniel Tannenbaum, as well as participants at Northwestern Applied Micro Lunch, the 2022 American Society of Health Economists Annual Conference, and the 2021 Association for Public Policy and Management Annual Conference. I thank Frank Limbrock and the Office of Statewide Health Planning and Development for data assistance. Financial support for this research is from the National Science Foundation Graduate Research Fellowship under Grant NSF DGE-1842165, and from the Susan Schmidt Bies Award.

I am so grateful for the support of my friends and fellow graduate students, particularly my office-mates, Eilidh Geddes and Udayan Vaidya.

Finally, this endeavor would not have been possible without my family. I am grateful for their encouragement along the way.

And to my husband, Justin Holz. Thank you for your never-ending support.

Dedication

To my family.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Information Disclosure and Patient Demand

1.1. Introduction

Consumers respond to information disclosure programs in many contexts, including education, charity, and healthcare (Allende et al., 2019; Dranove and Jin, 2010; Mayo, 2021). These disclosure programs are essential in the healthcare context, where spending is high and physicians vary in ways that impact spending and health outcomes Currie et al. (2016); Fadlon and Van Parys (2020). Further, outside of disclosure reports, patients do not have sufficient information to assess prospective physicians, especially on characteristics that they may not be able to evaluate on their own or learn from their peers. Disclosure reports aim to solve this problem by disclosing relevant cost and quality information about physicians. While many inputs impact the success of such programs, their efficacy in part depends on whether the patients who benefit the most or generate the most savings from seeing these physicians are also those who are most likely to respond to the disclosure program. This is especially true in contexts where capacity is constrained because the higher-rated physicians cannot treat all patients. If older patients have more to gain from seeing higher-quality, lower cost physicians, then disclosure policies could lead to larger cost savings by targeting information toward those patients.

I explore these heterogeneous effects in the context of a private health insurer's physician information disclosure program. The insurer creates quality and cost scores, which are continuous measures, based on the results of statistical tests applied to the physician compliance with quality measures and costs covered by both the insurer and the patients as compared to target benchmarks specific to the physician's case mix, specialty and geographic location. When there is sufficient data to evaluate performance, the insurer then uses the scores to determine if the physician meets the quality and/or cost criteria and then assigns physicians to one of three possible statuses: "lower-quality" (those who do not meet the quality criteria), "higher-quality" (those who meet the quality criteria, but not the cost criteria), and "higher-quality, lower-cost" (those who meet both quality and cost criteria).

The threshold rule facilitates a regression discontinuity (RD) design[1] that makes three comparisons. First, I compare "lower quality" physicians to "higher quality, lower cost" physicians by focusing on those who meet cost criteria and studying outcomes in the neighborhood of the quality threshold. These lower cost physicians who meet quality criteria and exceed the quality threshold are given "higher quality, lower cost" status. My findings show large gains in new patient volume for physicians in the neighborhood of the quality threshold with "higher quality, lower cost" status compared to "lower quality": they have 38% more new patient visits due to their status. The second comparison looks at physicians who do not meet the cost criteria and are given a "higher quality" (but not lower cost) status if they meet the quality criteria and are given "lower quality" status if they do not meet the criteria. These non–" lower cost" physicians experience low gains in the number of new patients in passing the quality threshold. Finally, I compare "higher

[1]The use of a regression discontinuity design is common in the quality disclosure field. For example, Anderson and Magruder (2012) and Chartock (2021) both use a regression discontinuity design to identify demand effects based on online information.

quality" to "higher quality, lower cost" physicians, for those in the neighborhood of the cost threshold. Among physicians who meet the quality criteria, those who exceed the cost threshold are given "higher quality, lower cost" status, and those who are below it are given "higher quality" status. The "higher quality, lower cost" status does not lead to statistically significant gains in the number of new patients when compared to the "higher quality" (but not lower cost) status. This is likely because the comparison of "higher quality, lower cost" to "higher quality" physicians is a comparison of the higher status to the middle status, whereas the comparison between "higher quality, lower cost" and "lower quality" physicians is a comparison of the higher status to the lower status, so potential effects in the latter comparison are larger.

The design is well-suited to identifying heterogeneity of the effect of the status on consumer demand because physician characteristics are continuous across the quality and cost thresholds. "Higher quality, lower cost" physicians are comparable to "lower-quality" physicians on every physician characteristic (such as age, gender, and geography) except the discrete status when their underlying quality score is close to the quality threshold. Any difference in the characteristics of the new patients that these physicians see represents heterogeneity based on patient demand responses to the disclosure program rather than physician characteristics.

I find that younger patients respond more to a physician's status than older patients: the average age of patients seen by "higher quality, lower cost" physicians is significantly lower (by 1.4 years) than the age of patients seen by "lower quality" physicians. This difference in the age of patients who see a "higher quality, lower cost" physician and those who do not can be interpreted as a result of heterogeneity in patient response. This is true so long as physicians on either side of the cutoff also behave similarly after scores are released. There is no empirical evidence that physicians change their behavior or game their score to achieve a "higher quality" (not low-cost) or "higher quality, lower cost" status, perhaps because scores are calculated using multiple years of claims (so averages are difficult to move), or because even though physicians know their own scores, they do not know others' scores or the distribution of scores.

The regression discontinuity estimates show how patients respond to the information program and reveal that younger patients respond more strongly to physician status, but the estimates do not shed light on which patients benefit the most in terms of spending or generate the most savings to insurers and patients collectively from seeing higher-quality, lower cost physicians. To study the heterogeneous effects of physician characteristics on patient outcomes, I complement the RD estimates with estimates from a two-way fixed effects "switchers" design, analyzing effects on total spending and preventable emergency room visits.

The two-way fixed effects strategy requires that the decision to switch physicians is exogenous to patient outcomes. I focus on individuals who change physicians the same year that their previous physician either moved or exited the dataset so that the decision to switch is plausibly exogenous. As a physician's cost score decreases, I find that their patients' spending decreases, on average, while an increase in a physician's quality score has the opposite effect (total spending slightly increases, on average). The mechanisms for these increases and decreases are different. Physicians with lower cost scores prescribe less expensive (but the same number of) services than physicians with higher cost scores. Conversely, higher-quality physicians prescribe more services compared to lower-quality physicians. While cost and quality scores impact various spending outcomes, there is no significant evidence that physicians who score higher on quality or lower on cost shift outcomes on emergency room utilization. Preventable emergencies do not significantly increase or decrease with a switch to a higher-quality or lower cost physicians.

Next, I turn to heterogeneity in patients' age by creating age-group subsets. I find that younger patients (aged 18–35) experience no significant gains or losses in total spending due to their physician's quality or cost score, while middle-aged patients (aged 36–55) experience larger impacts of both cost and quality. A back-of-the-envelope style analysis using estimates from the "switcher" model suggests that middle-aged patients are expected to spend about $30 fewer dollars annually (about 4% of average annual spending for this age group) than they would have otherwise when they switch from a "lower-quality" to a "higher quality, lower cost" physician. When the average younger patient makes the same switch, they experience a small and not statistically significant gain in annual spending of $0.28 (about 0.4% of average annual total spending for younger patients).

Younger patients are more responsive to physician status, but older, middle-aged patients experience the highest gains in savings from switching to "higher quality, lower cost" physicians. These results suggest that improved information targeting or larger financial incentives for middle-aged patients to respond to the physician's status may allow for larger net savings in healthcare spending without negative impacts on health, as measured by preventable emergency room visits.

This study contributes to the literature on quality disclosure (Dranove and Jin, 2010) by showing that privately insured patients respond to information about primary care physicians. The literature focuses primarily on how patients respond to ratings about hospitals and facilities (Bundorf et al., 2009) health plans (Wedig and Tai-Seale, 2002), and cardiac surgery (Yoon, 2020) or on the implications of providers adjusting their behavior in the presence of these programs (Dranove et al., 2003; Vatter, 2021). I instead focus on the effects of insurer-provided quality and cost scores on demand for primary care physicians. It is important to understand responses to information about primary care physicians because primary care physicians serve as care coordinators within the healthcare system by making referrals and providing essential preventative care services. The most closely related research by Chartock (2021) shows that patient satisfaction information impacts consumer choice for similar physicians. My study examines the effects of objective measures of quality—mainly adherence to clinical practice guidelines—that patients may respond to differently than measures of patient satisfaction.

This study also contributes in two ways to the quality disclosure literature's exploration of heterogeneity in disclosure effects. First, I consider several possible explanations for the age heterogeneity. This extends the research that quantifies similar patterns in heterogeneity. Beaulieu (2002) finds that younger consumers are more sensitive to price and less sensitive to physician networks than older patients, and Chartock (2021) finds that younger and healthier patients respond more to patient satisfaction scores. Second, my research provides empirical evidence that motivates learning about heterogeneity. The fact that the quality and cost scores of their physicians affects middle-aged patients more than younger patients gives researchers a reason to learn about heterogeneity in responses to information.

Finally, my study contributes to the literature that breaks down patient outcomes into patient- and physician-driven components. My results add more evidence to the literature, which finds that physician characteristics can impact patient outcomes[2]. Whereas researchers have generally focused on Medicare patients (Sabety, 2021; Fadlon and Van Parys, 2020; Kwok, 2019), this study focuses on the younger, privately insured population. My results showing that information disclosure about physicians affects middle-aged patients more than younger patients indicates the importance of examining the middle-aged population. A study focusing only on elderly Medicare patients may lead to different conclusions as results vary by age.

The paper proceeds as follows. Section 1.2 provides background information on the insurer's quality and cost disclosure program. In Sections 1.3 and 1.4, I introduce the study's data and empirical design. Section 1.5 presents results regarding the success of the program in influencing patients' selection of "higher quality, lower cost" physicians, and in Section 1.6, I explore the heterogeneity of the response by patient characteristics. Sections 1.7-1.9 discuss the data, empirical strategy, and results of the switcher analysis, which identifies the impacts of quality and cost on patient outcomes both on average and by age. Section 1.10 discusses limitations, and section 1.11 concludes.

[2]A large literature focuses on heterogeneity in provider practice styles and the impacts of provider-specific practice styles on patient outcomes. Chan et al. (2022), Currie and MacLeod (2017), Fletcher et al. (2014), Molitor (2018), Schnell (2017), Schnell and Currie (2018), and Van Parys (2016) establish that providers' characteristics vary, and they explore explanations for heterogeneous practice styles, such as training, altruism, skill level, or environmental factors. Currie and Zhang (2021), Dahlstrand (2021), Doyle Jr et al. (2010), Gowrisankaran et al. (2022) use random or quasi-random assignments of physicians to patients to understand the impacts of physician practice styles on patient outcomes, while Grytten and Sørensen (2003) and Simeonova et al. (2022) use fixed effects designs to separate patient-specific from physician-specific impacts.

1.2. Context: The Information Disclosure Program

I study the decisions of privately insured patients who choose primary care physicians in the context of an information disclosure program run by their insurance company.[3] Each year, the insurance company (hereinafter referred to as "Insurer X") uses medical insurance claims to construct binary indicators of cost and quality for physicians and discloses the indicators publicly online so that patients can use the information to choose a physician—someone who provides higher-quality, lower cost care. Insurer X's online physician directory shows each physician's status. When patients search for a physician in their insurance network using the online directory, they will see a list of physician profiles with status information. A physician is indicated a higher quality or higher quality and lower cost based predominantly on their quality and cost scores.

A physician who meets the quality criteria but not cost criteria is referred to as a "higher quality" physician. Physicians who have not met the quality criteria have no indicator and are referred to as lower quality. The information disclosure program began in the early 2000's, and each physician's status is updated each year based on claims from the prior two to three years. I study the 2019 program year because 2019 was the earliest year that program data were available, and 2020 data were likely impacted by the COVID-19 pandemic.

[3]This chapter discusses quality and cost metrics of physicians. The fact that a physician doesn't meet program quality and cost criteria doesn't mean the physician doesn't provide quality health care services. All physicians in the insurer's network have met certain minimum credentialing requirements. I use data provided by a large national health insurance company, which operates the disclosure program in 46 states in the United States. The data provider approved data questions before I conducted my research and it reviewed the draft of the paper, but it did not actively participate in the research process.

Insurer X uses standardized measures from the National Quality Forum and other sources[4] to calculate a physician's quality status. Using claims information, the company assigns each physician three numbers that make up the score. First, it assigns the number of "opportunities" the physician had to practice in accordance with the measures based on the physicians' attributed patient case mix. The number of opportunities is generally equivalent to the number of unique patient-disease treatment combinations attributed to the physician. For example, a patient with diabetes might count for three opportunities: the opportunity to perform a blood sugar control test, the opportunity to prescribe appropriate blood sugar control medication and the opportunity to check feet for abnormalities. The physician's compliance with the opportunities presented by each patient is compared to the median compliance rate, their "target benchmark" (at the national level, specific to the condition/procedure, physician specialty and population treated) for each opportunity. A statistical test is used to determine whether the physician's compliance is significantly less than the target benchmark. The physician meets the quality criteria when the physician's performance is not statistically less than the target benchmark. If it is, the physician's status is "lower quality". More specifically, for each physician, χ is calculated as shown in Equation 1.1:

$$\chi = \begin{cases} -1 \cdot \left[\frac{(P(c)-B(c))^2}{B(c)} + \frac{(P(n)-B(n))^2}{B(n)}\right] \text{ if } P(c) < B(c), \\ \left[\frac{(P(c)-B(c))^2}{B(c)} + \frac{(P(n)-B(n))^2}{B(n)}\right] \text{ if } P(c) \geq B(c) \end{cases} \tag{1.1}$$

[4]Most guidelines are from the National Quality Forum. Other sources include the National Committee for Quality Assurance, the Center for Medicare and Medicaid Services, and the Pharmacy Quality Alliance.

where c refers to the number of compliant measures, B to the target benchmark, and n is the number of non-compliant measures assigned to the physician, P. If χ is above a given cutoff, then the physician is considered a "higher quality" physician. I refer to those who missed the cutoff as "lower quality" physicians or "physicians who did not meet quality criteria."

Insurer X also uses information from claims to assign each physician's cost status. Insurer X creates a measure of the total cost of a patient over a single calendar year and attributes that cost to the physician with the most significant involvement in the patient's care. The total cost measure aggregates costs for conditions and procedures relevant to the physician's specialty in determining differences in physician performance. Qualifying payments are determined by specialty. For example, costs associated with the diagnosis of allergic rhinitis are included in the cost measure for allergists. The total cost variable includes both the patient's and the insurer's portion of the payment. The total cost is then risk-adjusted using the patient's risk score based on the patient's underlying clinical characteristics. Like the quality calculation, a benchmark score is then calculated based on the 75th percentile cost level, and a statistical test is used to determine whether the adjusted costs attributed to the physician were significantly lower than the benchmark score. I will call physicians "lower cost" if their cost score lies below the threshold, and "non–lower cost" if their score lies above the threshold.

Each physician's online profile contains other information. The name of the physician is included, as well as the physician's main address and the distance from the patient to the physician's office. If the physician has a profile on Healthgrades.com[5] (a popular

[5]Healthgrades scores are available publicly at https://www.healthgrades.com

physician rating website, where ratings are based on patient satisfaction), the average star rating is shown. A patient can sort physicians based on gender, distance, or quality/cost status; however, the default sorting is based in part by the physician's status.

1.3. Data – Physician Regression Discontinuity

I use data from the information disclosure program in 2019 to measure each physician's underlying quality and cost score as well as their discrete publicly available quality and cost status. I use data on the underlying scores and on the resulting status of each physician. Table 1.1 displays the number and proportion of physicians with each status in the bottom two rows.

When exploring patient response to physician status, the main outcome variable is the number of new patients seen by each physician each month and each quarter. To create this variable, I begin with a dataset for each month of claims with patient and physician identifiers and procedure (CPT) codes. If the procedure code used corresponds to a new patient visit,[6] then I label the patient as a new patient. The new patient indicator only counts first visits: if the new patient returns to the physician in the same month, only one new patient visit is accounted for. For most specifications, I create a cross-sectional dataset by aggregating to the physician level, including new patient visits for the three months after the new quality scores were released in 2019.[7] For specifications that rely on variation over time, I use physician-month–level data. Table 1.1 shows the number of total and new patients by status in the cross-section. "Higher quality, lower cost" physicians tend to see slightly more new and returning patients.

[6]Appendix Table A.2 lists the CPT codes used to designate new patients.
[7]Program data prior to 2019 are unavailable, so I focus on the 2019 update to avoid disruptions caused by the COVID-19 pandemic.

	Lower Quality	Higher Quality, Not Lower Cost	Higher Quality, Lower Cost	All
Total Patients	36.56	38.46	40.14	38.90
	(60.78)	(164.56)	(51.70)	(121.48)
New Patients	4.30	4.00	4.38	4.20
	(10.60)	(7.82)	(8.06)	(8.30)
Older New Patients	2.00	2.08	2.10	2.08
	(4.80)	(3.72)	(3.58)	(3.82)
Younger New Patients	2.32	1.94	2.28	2.12
	(6.50)	(5.08)	(5.30)	(5.36)
Chronically Ill New Patients	0.96	0.92	1.02	0.96
	(2.24)	(1.84)	(1.94)	(1.92)
Non-Chronically Ill New Patients	3.36	3.08	3.38	3.22
	(8.74)	(6.44)	(6.60)	(6.82)
Observations	10616.00	44347.00	35519.00	90482.00
Proportion of Total	0.12	0.50	0.40	1.00

Table 1.1. Summary Statistics

Notes: This table reports summary statistics from claims with a date between zero and three months after disclosure, broken down by provider status. Older new patients are those patients aged 40 and older who saw their doctor for the first time, while younger new patients are those aged 18-39 who saw their doctor for the first time. Chronically ill patients are those who have a diagnosis of any chronic illness from an office visit over the past year. Proportions as displayed are rounded, and rounding error explains why the proportions do not sum to one. Standard errors are shown in parentheses.

I also create variables for the number of new patients with various other characteristics, such as the number of new patients who are younger than 40, or the number of new, chronically ill patients. While some characteristics, such as age and gender, are available directly in the medical claims, the chronic illness indicator must be constructed. I use the methodology from Gruber and McKnight (2016) to construct chronic illness indicators. For each patient month, I determine whether, in the prior year, the patient was seen in an office setting for a diagnosis of one of six common chronic illnesses[8] If they were, then the patient is designated as chronically ill for that month.

To explore mechanisms of heterogeneous responses to physician status, I create measures of situations that may impact how patients search: whether the patient or physician recently moved, whether the patient is new to the insurance company, and how much the patient paid out of pocket for the visit. First, I denote whether a patient or their physician has recently moved based on whether the first two digits of the patient's zip code or the physician's zip code matches the first two digits of the zip code the patient had during their most recent prior visit to a physician of the same specialty. Next, I create an indicator for whether the patient is new to the insurance company based on their enrollment dates. If the patient was not insured during 2018 but was insured during 2019, the year studied, the patient is considered "new to the insurer." Finally, I explore the amount each patient paid out of pocket for the visit. Although a patient may not know the exact amount they will be billed for a given visit ahead of time, they likely have a general idea of what their visit will cost based on whether they have already met their deductible or out-of-pocket maximum, and whether the type of care is fully covered (for example, an annual wellness check is fully covered). Patients may search for lower cost physicians based on their expectations of how expensive their care will be.

Finally, I merge two additional physician covariates: average ratings from Healthgrades.com, and physician medical school graduation year. The Healthgrades score is merged based on name and location of the physician. About 12% of physicians within the bandwidth successfully merged to Healthgrades data. Although the match rate is quite

[8]I follow Gruber and McKnight (2016) in identifying the following chronic illnesses: diabetes, asthma, arthritis, affective disorders, and gastritis.

low, matching exactly on name and location ensures that matches are more likely to be accurate. Graduation year can be merged directly on physician identification numbers (NPIs); thus, the match rate is higher at 41%. However, graduation year is only available for physicians who also treat Medicare patients. Since merge rates are low, most analyses omit these variables to allow for the analysis of a larger dataset.

1.4. Empirical Strategy: Identifying the Physician Status Effect

I use a regression discontinuity design to determine the impact of quality status on the number of patients the physician saw over the three months after designations were updated. Physicians are given status based predominately on quality and cost scores. I make three comparisons, first identifying the effect of having higher quality, lower cost" status versus no "lower quality status by looking within lower-cost physicians and comparing those around the quality threshold. I next compare higher quality to lower quality by looking within non-lower cost physicians and comparing those with scores around the quality threshold. Finally, I compare higher quality, lower cost to higher quality by looking within higher-quality physicians and comparing those with scores around the cost threshold.

I analyze these three comparisons by creating a subset of the sample of physicians in different ways. First, to determine the impact of "higher quality, lower cost" compared to "lower quality", I create a subset of lower cost physicians who are eligible for "higher quality, lower cost" status because they have low enough cost scores. Comparing lower cost physicians just to the left versus just to the right of the quality cutoff identifies the impact of "higher quality, lower cost" status (right of cutoff) relative to being "lower quality" (left of cutoff). Likewise, comparing non-lower cost physician just to the left versus just to the right of the quality cutoff identifies the impact of being classified as "higher quality" physician relative to being "lower quality" because non-lower cost physicians are given the "higher quality" status rather than the "higher quality, lower cost" status upon passing the quality threshold (right of cutoff). Finally, creating a subset of higher-quality physicians, I compare physicians just to the left versus those just to the right of the cost threshold. Those to the left are "higher-quality, lower-cost" physicians (as a lower cost score is better in this case), while those to the right are higher-quality, non–lower cost "higher quality" physicians.

Each status is displayed on physician directories. Thus, the first natural experiment that compares "higher quality, lower cost" physicians to "lower quality" physicians compares physicians with higher quality, lower cost to those with no indicator. The second natural experiment, comparing "higher quality" to "lower quality" physicians, compares physicians with higher quality to those with no indicator. The third natural experiment compares physicians with higher quality, lower cost to those with only higher quality.

To measure the difference in outcomes between physicians whose quality or cost scores were just above versus just below the threshold, I estimate an ordinary least squares regression within a small bandwidth around the cutoff, within the relevant subset (e.g., only lower cost physicians). Specifically, I regress

$$NewPatients_j = \alpha + \delta_1 X_j + \beta 1\{X_j > 0\} + \delta_2 X_j 1\{X_j > 0\} + \epsilon_j \ , \ X_j \in [-10, 5] \tag{1.2}$$

where $NewPatients_j$ measures the number of new patients seen by physician j within the first three months of designation publication. X_j measures the continuous, underlying quality score, which is standardized so that the cutoff is zero. β measures the impact of passing the quality threshold and being classified as a "higher-quality" physician. The process is repeated with X_j as the standardized cost score to estimate the impact of being classified as a lower cost physician (comparing "higher quality, lower cost" to "higher quality" (not low cost) physicians).

To visually assess whether discontinuities exist, I plot outcomes over bins of the running variable: either the cost or quality score. I partition the running variable into 20 quantiles on either side of the discontinuity within the bandwidth and display the average outcome within each bin. I also plot the regression discontinuity predicted values (lines on either side of the cutoff) to help visualize the exact size of the discontinuity. I use a bandwidth that extends to 10 points below zero on the left of the cutoff and 5 points above zero to the right of the cutoff for all quality regression discontinuity plots, and a bandwidth extending to 3 points on either side of the cutoff for plots where the cost score is the running variable. These bandwidths were chosen because they are close to the mean squared error optimal bandwidths (Calonico et al., 2014) for multiple specifications. Since optimal bandwidths vary between specifications, using fixed bandwidths of [–10,5] and [–3,3] allows for the bandwidths to stay the same across all specifications. Because data are sparser on the left than the right of the quality threshold, the mean squared error optimal bandwidths for the analyses around the quality thresholds are calculated allowing for different bandwidth sizes on the left and right. Results from iterating bandwidths are shown in Appendix Figure A.12.

1.4.1. Validity checks and evidence against gaming

Interpreting β as a causal effect requires that physicians are not able to engage in "gaming" by selecting which side of the cutoff they are on. It is unlikely that physicians game in this way. Physicians are given detailed information of their performance on various measures; however, they are not told how performance on each measure maps to their quality score. To back out the mapping, a physician would need to know the national median compliance rate for each measure, the exact number of patients he treated with each condition, and the algebraic formula to calculate the resulting quality and cost scores. Although this information is available to physicians in principle, it is unlikely they would spend the time to calculate exactly how much additional care to give to specific patients. Finally, for a physician to respond to the program by changing behavior, the physician would have to run another diagnostic test or make another pharmaceutical prescription for a patient for whom he would previously have not done this test or prescribed medication. It is difficult to imagine a physician who knew she was overlooking these tests for specific patients but did not decide to correct this behavior until after she was notified of her lower-quality status. If the physician had known of her noncompliance all along, why would she not have corrected this behavior earlier?

The best evidence against physician gaming, however, is empirical. Figure 1.1 depicts density plots for all three natural experiments (in the spirit of McCrary (2008). If physicians can manipulate their scores to improve their likelihood of being classified as "higher quality" or "higher quality, lower cost" physicians, I would expect to see missing mass just to the left of the cutoff and additional mass just to the right of the cutoff, but no such patterns are detected in the data. I also run the density tests outlined by Calonico et al. (2014) and find no statistically significant discontinuities in the density of physicians around the cost or quality thresholds. Appendix Figure A.4 displays the results of these density tests.

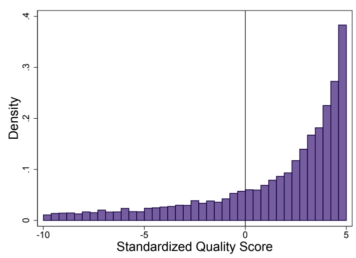

(a) Lower-cost Physicians Around Quality Cutoff

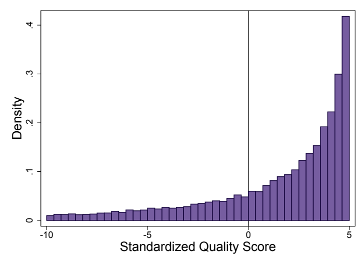

(b) Not Lower-Cost Physicians Around Quality Cutoff

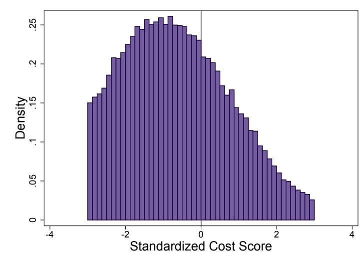

(c) Higher Quality Physicians Around Cost Cutoff

Figure 1.1. Density Plots
Notes: This figure displays three density plots. Panel A displays the density of underlying quality scores for lower-cost physicians. Panel B displays the same for not lower-cost physicians, while panel C displays a histogram of underlying cost scores for higher-quality physicians. There is no visual evidence of bunching around any cutoffs.

If physicians are as good as randomly assigned around the cutoff, I would also expect that physician characteristics are continuous at the cutoff. To assess continuity around the quality cutoff, I predict the number of new patients seen by each physician as a function of physician gender and zip code fixed effects on the cross-section of data from the three months prior to the 2019 program update. I then estimate Equation 1.2 with the predicted number of new patients as the outcome, shown in Figure 1.7, for each natural experiment. I find no significant impact or visual evidence of a discontinuity. Appendix Table A.3 shows the regressions used to predict new patients.[9]

[9]Graduation year is only available for physicians who see Medicare patients, and Average Healthgrades Score is only available for physicians who have a Healthgrades.com profile and were successfully matched to the claims dataset on name and location. These variables are omitted from the prediction of new patients because they are only populated for a subset of variables; however, the number of predicted new patients using all five variables is also smooth around the cutoff.

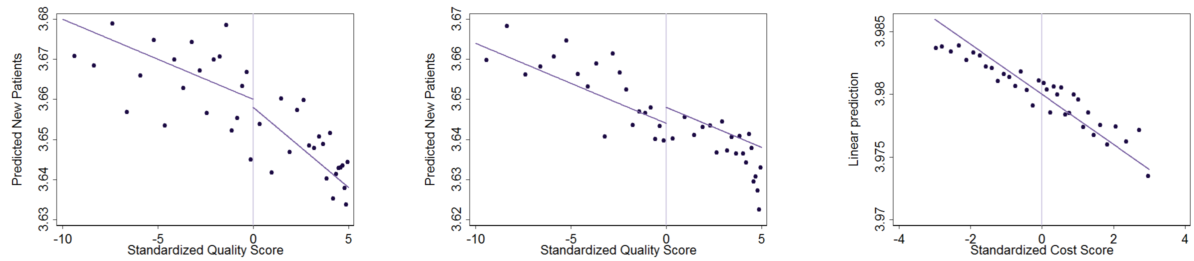

(a) Lower Quality vs. Higher Quality, Lower Cost

(b) Lower Quality vs. Higher Quality, Not Lower Cost

(c) Higher Quality, Not Lower Cost vs. Higher Quality, Lower Cost

Figure 1.2. Validity Test: Impacts on Predicted New Patients
Notes: This figure displays three RD plots. All three display binned scatter plots of predicted number of new patients based solely on provider characteristics. If RD assumptions are satisfied, predicted new patients should be smooth around the cutoff. Panel A displays predicted patients around the quality threshold for lower cost physicians, who were assigned "lower quality" status if they missed the quality threshold, and "higher quality, lower cost" status if they achieved it. Panel B displays the same for not lower cost physicians, who were assigned "higher quality" (not lower cost) status if they passed the threshold, and "lower quality" status if they missed it. Panel C displays predicted values of new patients around the cost threshold for higher quality physicians. Those to the right of the cutoff were assigned "higher quality" (not lower cost) status, while those to the left were assigned "higher quality, lower cost" status. None of the figures show evidence of physician gaming or manipulating the threshold.

I also estimate the regression in Equation 1.2 with each physician characteristic as the outcome and report results in Table 1.2. In Table 1.2, I also measure impacts on graduation year and average patient review scores from Healthgrades.com, which are only available for a subset of physicians. I find no evidence of imbalance on these characteristics.

1.4.2. Identifying patient versus physician responses using differential timing

I use a difference-in-discontinuity approach to identify whether physicians or patients are responsible for any differences in outcomes. Although intuitively patients would be

	(1) Female	(2) Cost Score	(3) Graduation Year	(4) Healthgrades Score
Panel A: Lower-Cost Doctors				
Above Quality Cutoff	0.0130	0.0228	-0.903	-0.142
	(0.0211)	(0.0472)	(0.780)	(0.129)
Observations	8743	8743	4305	843
	(1)	(2)	(3)	(4)
Panel B: Not Lower Cost Doctors				
Above Quality Cutoff	-0.0172	-0.0767*	-0.733	0.0590
	(0.0151)	(0.0411)	(0.548)	(0.106)
Observations	19403	19403	9027	1999
	(1)	(2)	(3)	(4)
Panel C: Cost Threshold - Higher Quality Doctors				
Below cost cutoff	0.00484	0.140	-0.0526	0.0315
	(0.00714)	(0.0919)	(0.242)	(0.0390)
Observations	67644	67644	30545	7707

Table 1.2. Covariate Continuity

Notes: This table reports results of the regress ion in equation 2 on physician characteristics as outcomes.

more likely to respond to the information, in principle it is possible that physicians could respond as well, perhaps by turning away new patients or advertising more heavily. The approach leverages disclosure timing: the insurance company notified physicians of their status two months before the company updated the physician statuses online. If effects materialize before patients could have seen updated statuses but after physicians were notified, then results could be attributed to physician behavior, whereas if results do not materialize until after physicians were notified of their score, the impacts were likely due to patient responses.

I estimate the following regression equation, adapted from Grembi et al. (2016) to break down effects year-by-year in an event study–style framework:

(1.3)

$$NewPatients_{jt} = \sum_{k=-5}^{8} 1\{t = k\}[\beta_k 1\{X_j > 0\} + \gamma_{1k} X_j + \gamma_{2k} X_j 1\{X_j > 0\}] + \tau_t + \alpha_j + \nu_{jt}$$

where X_j represents the running variable—underlying quality or cost score—and τ_t and α_j are month and physician fixed effects. β_k estimates the differential impact of having passed the threshold in year k relative to the effect in the month before the scores were given to physicians. If β_ks are close to zero and not significant during the time when physicians knew their updated status but patients did not, then we can infer that physicians did not respond immediately to the updated information. If effects materialize only after scores were available publicly to patients, then one can infer that patients did respond and physicians did not under two assumptions. First, one must assume that impacts on physicians were constant over time so that the new information did not take more than two months to change the behavior of physicians. Second, one must assume that there are no complementarities between physicians and patients both knowing information that could impact physician behavior (physicians cannot have waited for the public disclosure date to change their behavior). See Grembi et al. (2016) for further discussion of these assumptions.

1.4.3. Patient heterogeneity

Measuring heterogeneous responses to a physician's designation is not straightforward because the treatment is at the physician level. To determine whether certain groups of patients respond more than others, I use two approaches. First, I use patient-physician–level data from the three months after the 2019 program update and explore differences in characteristics of patients seen by various groups of physicians. Because physician characteristics are held constant by the identifying assumption of the regression discontinuity design, any discontinuity in patient characteristics can be attributed to either patients responding differently based on their characteristics or by physicians purposefully targeting specific types of patients. I estimate Equation 1.2 with the average age and chronic illness status of new patients as outcomes, regressing at the patient level to avoid any aggregation bias that might result from first collapsing to the physician level, and then running the regression. I cluster standard errors at the physician level to account for any within-physician correlation in chronic illness status and age.[10] I additionally estimate Equation 1.2 on the number of new patients who are younger than 40 years of age and on the number of new patients who are older than 40 years of age and likewise for chronic illness status.[11] To determine whether changes are due to physician or patient changes, I also estimate Equation 1.3, exploring dynamic effects, on patient characteristics. These complimentary analyses provide another way of examining heterogeneity in responses to the physician's designation.

[10]Running regression 1.2 "at the patient level" means estimating the following regression: $Y_{ij} = \alpha + \delta_1 X_j + \beta_1 1\{X_j > 0\} + \delta_2 X_j 1\{X_j > 0\} + \epsilon_{ij}, X \in [-10, 5]$, where Y_{ij} is the patient-level characteristic such as age or chronic illness status.
[11]Privately insured adults are usually under age 65, so 40 years is roughly between ages 18–65.

1.5. Physician Status Effect Results

1.5.1. The Physician Status Effect

I begin by assessing the first stage. I break physicians into cost groups based on whether the cost score exceeds the cost threshold. For lower cost- physicians, passing the quality threshold constitutes a sharp regression discontinuity (RD): everyone whose quality score exceeds the cutoff has "higher quality, lower cost" status. For physicians who are not "lower cost," passing the quality threshold should lead to "higher quality" status. However, physicians who do not have a sufficient number of attributed measures may meet the criteria for "higher quality, lower cost" status based on an affiliated medial group evaluation result for their specialty within the same geographic area.[12] This means that some discontinuities are fuzzy because some physicians whose cost scores would have been too high achieve "higher quality, lower cost" status instead of "higher quality" status upon passing the quality threshold because of their affiliated group's evaluation results. Appendix Figure A.5 displays these first-stage plots in Panels (a) and (b).

I next turn to the main results. I break down physicians into two categories based on their cost status. Plotting new patients over physician quality in Figure 1.3 reveals a discontinuity around the quality threshold for primary care physicians[13] but only for the lower cost physicians who achieve "higher quality, lower cost" status upon passing the quality threshold. A line is plotted on either side of the threshold, within the bandwidth. Passing the quality threshold leads to seeing more new patients for lower cost primary care

[12]Group quality evaluation results are based on services provided to patients enrolled in Insurer X commercial fee-for-service, Medicare Advantage, and community health plans. Group cost evaluation results are based on commercial treatment sets only. Medical group evaluation benefits physicians with an affiliated medical group who share practice patterns and care protocols.
[13]Impacts for specialists are displayed in Appendix Figure A.6.

physicians. For non–lower cost physicians impacts are much smaller and not statistically significant.

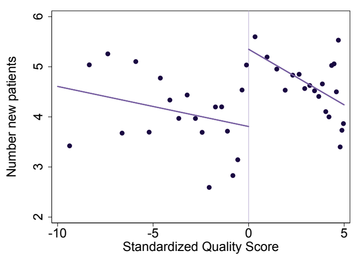

(a) Lower Quality vs. Higher Quality, Lower Cost

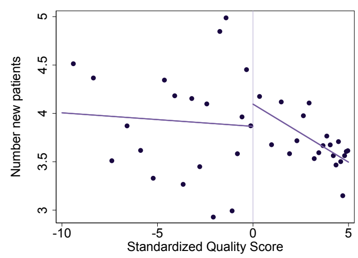

(b) Lower Quality vs. Higher Quality, Not Lower Cost

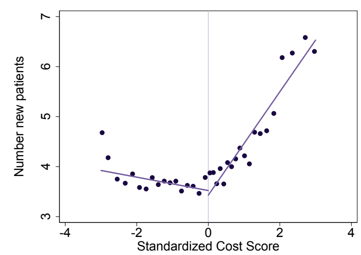

(c) Higher Quality, Not Lower Cost vs. Higher Quality, Lower Cost

Figure 1.3. These figures plot the number of new patients seen by providers over the three months following status updates. Outcomes are averaged in each bin for twenty quantile-spaced bins around the cutoff. Panel A shows the impact on new patients for lower-cost primary care physicians. These providers achieved "lower quality" status if below the cutoff, and "higher quality, lower cost" status if above. Panel B shows the impact on new patients for not lower-cost physicians (comparing "higher quality" (not lower cost) to "lower quality"), and Panel C shows the impact on new patients for higher-quality physicians based on whether the physician was below the cost threshold (comparing "higher quality, lower cost" status to the left of the cutoff to "higher quality"(not lower cost) status to the right of the cutoff).

Panel (c) of Figure 1.3 displays results for the regression discontinuity estimation around the cost threshold, comparing higher quality physicians with higher quality lower cost physicians. There is no visual discontinuity for either specialists or primary care physicians, and the regression equation picks up a null result (all regression results are displayed in Table 1.3). The null result could come from the fact that "higher quality, lower cost" status is the higher of the three statuses, and when it is compared to the middle status, "higher quality" (not lower cost), the effects are not large enough to be detected. Or, the result could be because the cost RD is fuzzy so the comparison across the cost threshold includes some physicians on both sides of the threshold with a "higher quality, lower cost" status, which would attenuate RD results. One way to determine whether the fuzzy nature of the RD is attenuating results is to run an instrumental variables (IV) regression, where passing the threshold is used as an instrument for lower cost status. The regression results show that having lower cost status leads to an insignificant increase in the number of new patients seen by physicians who barely passed the cost threshold. Appendix Table A.4 displays the results of this regression. The IV results are not statistically significant, so it is unlikely that the fuzziness of the discontinuity is hiding a significant impact. However, it is possible that the effects of higher quality increases are too small to be detected in the sample, whereas the large effect comparing higher quality, lower cost versus lower quality is detectable by the sample size. Further discussion of the differences between achieving a "higher quality" and achieving "higher quality, lower cost" is in Appendix A.1.

	(1) Number new patients	(2) Number new patients	(3) Number new patients	(4) Number new patients
Above Quality Cutoff=1	0.642***	1.543***	0.230	
	(0.239)	(0.437)	(0.286)	
Below Cost Cutoff=1				-0.0921
				(0.110)
Heterogeneity Variable	Primary Care	Primary Care	Primary Care	Primary Care
Cost or Quality Status	All	Lower Cost	Not Lower Cost	Higher Quality
R^2	0.000592	0.00187	0.000553	0.00602
Observations	28146	8743	19403	67644

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3. Main Effects for Primary Care Providers

Notes: This table displays main effects for primary care providers. Column 1 displays pooled results over cost status, while columns 2 and 3 show results for lower and non-lower-cost physicians separately. Column 4 displays the impact of passing the cost threshold for the higher quality providers who met the quality criteria.

The fact that lower cost physicians drive the effects makes sense for multiple reasons. First, since sorting is in part based on status, achieving "higher quality, lower cost" status may increase the page rank of physicians substantially, possibly to the first page of search results. While "higher quality" (not low-cost) physicians also experience an increase in page rank, the increase is likely smaller, and the likelihood that the first page displays the physicians' profile is also smaller. The effects shown can be considered as the effect of being classified with a given status, which includes both the ranking effect and the information effect. Disentangling these two effects is beyond the scope of this study because data on page rank are not available.

1.5.2. Patient versus physician responses

In evaluating the cross-sectional results, one may wonder whether the increases in new patients are truly due to patient demand as opposed to physician behavior such as turning away new patients or changing other policies that may affect the number of new patients a physician sees. The timing of the information disclosure can be used to determine whether patients or physicians are responding as physicians were informed of their status about two months before the scores were publicly visible to patients. Figure 1.4 displays the results of the difference-in-discontinuity regression shown in Equation 1.3. Each β_k is plotted over time, and the two vertical purple lines show the timing of physician and patient disclosures. Results for lower cost primary care physicians are shown; because the impacts on non-lower cost physicians are limited, there is little concern of physician behavior impacting results for those specifications. For this reason, I focus only on lower cost primary care physicians for the remainder of the paper.

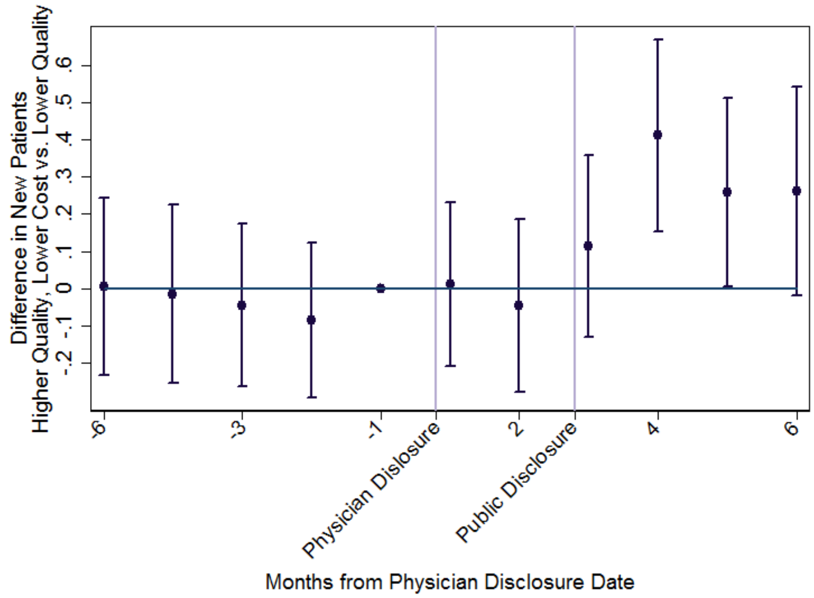

Figure 1.4. Difference in Discontinuities
Notes: This figure shows β_k for months leading up to and following disclosure of updated statuses to providers and patients for lower-cost primary care providers.

There is no effect (relative to the month before physician disclosure) of being classified as a "higher quality, lower cost" physician in the 2019 classification between February and June of 2019, before physicians or patients were notified of their designation. This also serves as a validity test, showing first that physicians who would just pass the quality threshold in the 2019 update were trending similarly in terms of new patients as physicians who would just fail to meet the quality threshold. If physicians were gaming the cutoff, one might instead expect to see some form of pre-trends in this figure. The fact that no pre-trends are observed gives credence to the assumption that physicians near the cutoff are comparable.

During the two months when physicians knew their status but patients did not, the effects stay around zero, whereas after new physician statuses were displayed at the end of

September, the effects began to materialize. By October 2019 (four months post-physician disclosure), the full effect was realized. The dynamic effects show that patient demand most likely drives the empirical results because no effects are seen until after patients were able to access updated physician statuses. On the one hand, one cannot rule out that physicians changed something about their capacity or likelihood of taking on new patients in a way that results in lagged effects. However, because the average wait time for a patient appointment for primary care physicians is 40 days (Penn et al., 2019), it is likely that if physicians influenced the number of new patients they saw, effects would have been visible at least by August.

Collectively, the results in this section show that patients respond to the information program and that effects are concentrated in "higher quality, lower cost" physicians who have higher page ranks. The results are likely due to patient behavior because the effects do not materialize until after patients have access to the information.

1.6. Heterogeneity

To determine whether specific types of patients respond more than others, I estimate Equation 1.2 at the patient level on patient characteristics as outcomes. If certain types of patients respond strongly to a physician's status, the average characteristics of patients treated by those physicians will change. Table 1.4 displays the results of regressions on average age, chronic illness status, and previous year's spending for new patients treated by lower-cost primary care physicians who barely met the quality threshold versus those who barely missed it. My findings show that meeting the quality threshold leads to a relative decrease in the average age of patients and a marginally significant relative decrease in chronically ill patients. There is no impact on the previous year's spending of the patients seen, meaning that those who had higher spending in the previous year did not respond more or less to physician status than those who had lower spending in the previous year.

	(1) Age	(2) Chronic Illness	(3) Last Year's Spending
Above Quality Cutoff	-1.407**	-0.0243*	-80.61
	(0.613)	(0.0127)	(170.7)
Observations	39176	39176	39176

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.4. Patient Heterogeneity

Notes: This table displays the regression discontinuity estimation with patient characteristics as outcomes for lower-cost primary care physicians. Column 1 displays the impact of the "higher quality, lower cost" status on the average age of new patients seen, and columns 2 and 3 report impacts on chronic illness status and the patient's prior year total spending (insurer + patient). The status itself does not change patient characteristics, instead, these impacts should be interpreted as resulting from different responses of patients to "higher quality, lower cost" status based on the patient's characteristics.

The average age and chronic illness status of patients could be decreasing for primary care physicians either because young patients respond more to physician status or because older patients respond negatively to physician status. To evaluate these possibilities, I display regression discontinuity plots in Figure 1.5 separately for new patients who are younger than 40 versus new patients who are 40 or older. I standardize the measure of new patients for each group as the base rate of new patients is different between groups, so changes in the standardized measure can be interpreted as proportional increases from the mean of each outcome. Figure 1.5 depicts that while both older and younger patients respond positively to the status of lower cost primary care physicians, younger patients respond more. Likewise, both chronically ill and healthy patients respond, but healthier patients respond slightly more than chronically ill patients respond to the lower cost status. Figures displaying the regression discontinuity plots for non-lower cost physicians are in Appendix Figure A.8.

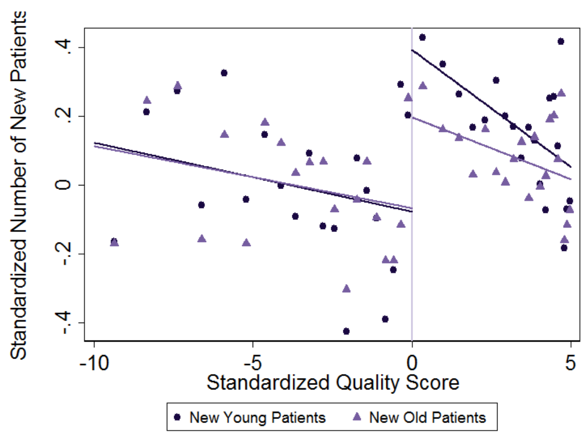

(a) Younger Versus Older Patients

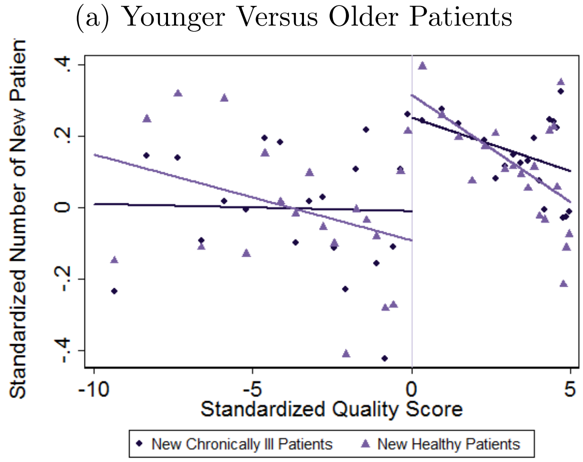

(b) Chronically Ill Versus Healthy Patients

Figure 1.5. These figures plot the number of new patients seen by providers over the three months following status updates for lower-cost primary care providers. New patients are broken down into two groups based on their age (above versus below 40 years of age) patients in panel A and into chronically ill versus healthy patients in panel B. Outcomes are averaged in each bin for twenty quantile-spaced bins around the cutoff.

To interpret these impacts as being due to patient demand rather than changes in physician behavior, I estimate Equation 1.3 at the patient level on patient outcomes. If patient characteristics change before patients can view updated scores, then the differential impacts may be due to physicians changing the types of patients they see rather than different patients responding differently to the information. Appendix Figure A.7 displays impacts on average age and chronic illness for lower cost primary care physicians over time. My findings show that impacts on average age are significantly lower during post-disclosure months. Impacts on average chronic illness status are less conclusive: there is some evidence of slight relative increases in the level of chronic illness; however, some pre-periods have significant increases as well. For age, responses are likely driven by patient behavior, but heterogeneity in chronic illness status does not hold up to this robustness test. Therefore, in the remainder of this paper, I focus only on how a patient's age impacts his or her response to the information program.

For further evidence of the pattern that younger patients respond more than older patients, in Appendix Figure A.13 panel (a) I display the raw data plot which shows the average age of patients seen by lower cost physicians around the quality cutoff, visually confirming the discontinuity estimated in Table 1.4. I also break down the impacts of passing the quality cutoff on new patients in six age categories in panel (b) to explore patterns over the age distribution. These exercises confirm that younger patients respond more than older patients do.

1.6.1. Explaining heterogeneity

Younger patients respond more to physicians' status than older patients do when they see primary care physicians. Is this due to the search behaviors specific to younger patients versus older patients, or do other characteristics that correlate with age and health impact search behavior? Understanding drivers of age-related heterogeneity is important so that policymakers can more effectively target information in future programs.

First, I evaluate one possible reason for the heterogeneity: that older patients are more attached to their physicians than younger patients are, which leads to reduced switching for older patients. This is true in the dataset: the average number of older new patients is smaller than the average age of younger new patients. However, if this difference explained the heterogeneity (i.e., that responses were proportionally the same but the base rates are lower for older patients), then the standardized number of new patients displayed in Figure 1.6, which nets out the average new patients for each group, would not display differences in effects for older versus younger new patients. Further, the impacts on average age displayed in Table 1.4 already adjust for different base rates because those base rates play a role in the average age of patients on both sides of the cutoff. If different base rates did explain results, then no impacts would be found. Because different base rates do not explain results, then something else must explain the heterogeneity.

Next, I evaluate whether three search-related characteristics explain the age and health-status related heterogeneity. While younger patients are more likely to have recently moved and are more likely to be newly insured (as shown in Figure 1.6), there is no substantial heterogeneity based on these characteristics, which one might have thought would impact the search process. Individuals who recently moved may have a smaller

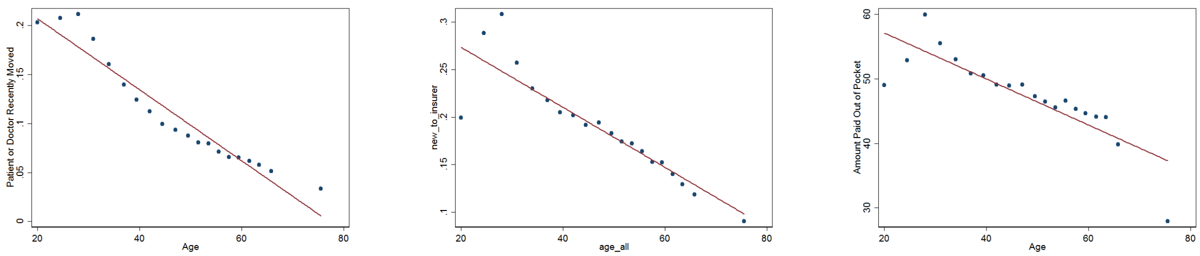

(a) Recently Moved (b) New to Insurer (c) Out of Pocket Payment

Figure 1.6. Binned Scatterplots of Age Against Search Characteristics
Notes: These figures display binned scatterplots for patient characteristics against age. Younger patients are more likely to have recently moved or for their physician to have recently moved, are more likely to be new to the insurer, and face higher out of pocket payments for their upcoming visit.

network to rely on when searching for a physician and may therefore have to rely more on website information. Individuals who are new to the insurer may have less exposure to the program and therefore may not know about the program or know how to search through the insurer's online website. Table 5 displays results of estimating Equation1.2 at the patient level on outcomes for whether the patient or their physician recently moved, whether they were new to the insurer in 2019, and the amount spent by the patient for lower cost primary care physicians. My findings show no significant increases or decreases in the proportion of patients who recently moved or who were new to the insurer; however, there is a significant increase in the amount paid by patients who saw lower cost physicians who just exceeded the quality score cutoff.

The increase in the amount paid by patients who saw "higher quality, lower cost" providers who just passed the quality cutoff suggests that financial considerations may explain heterogeneous effects by patient age. This result should not be interpreted as information about what happens to patient spending when they see a "higher quality, lower cost" physician because the "higher quality, lower cost" physicians who just passed

	(1) Patient or Doctor Moved	(2) New to Insurer	(3) Out of Pocket
Above Quality Cutoff	-0.00598	0.00173	10.44**
	(0.00934)	(0.0130)	(5.283)
Observations	39176	39176	39174

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5. Patient Search Mechanisms

Notes: This table displays the results of individual level regressions estimating equation 2 on characteristics related to patient search: whether the patient or physician recently moved, whether the patient is new to the insurer, and the patient's out of pocket payment.

the quality threshold are comparable in cost scores (and other underlying characteristics) to those who just missed the threshold. Instead, the impacts should be interpreted as identifying differences in the types of patients (those who were about to spend more versus less) who choose to see "higher quality, lower cost" versus "lower quality" physicians. As age is negatively correlated with out-of-pocket spending, it could be the case that young patients are more responsive to information about physician costliness than older patients because they have more to gain in terms of cost savings. This hypothesis is further supported by the fact that heterogeneity based on age is only present for lower cost physicians, whereas young and old patients respond similarly to the "higher quality" (not lower cost) status (see Appendix Figure A.8).

Finally, I consider whether access to the information via the internet explains the differential effects by patient age. Using data from the 2011 National Health Interview Survey (*National Health Interview Survey 2011 Data Release*, 2011), I plot the proportion of people in each age group who answered "Yes" to a question asking if the respondent had looked up health information on the internet in the past 12 months. The results, shown in Figure 1.7, suggest that access to the information could explain different responses by age. The pattern of internet search by age, where the youngest respondents were most likely to look up health information online and the oldest were least likely, matches almost exactly the pattern of response to the disclosure program.

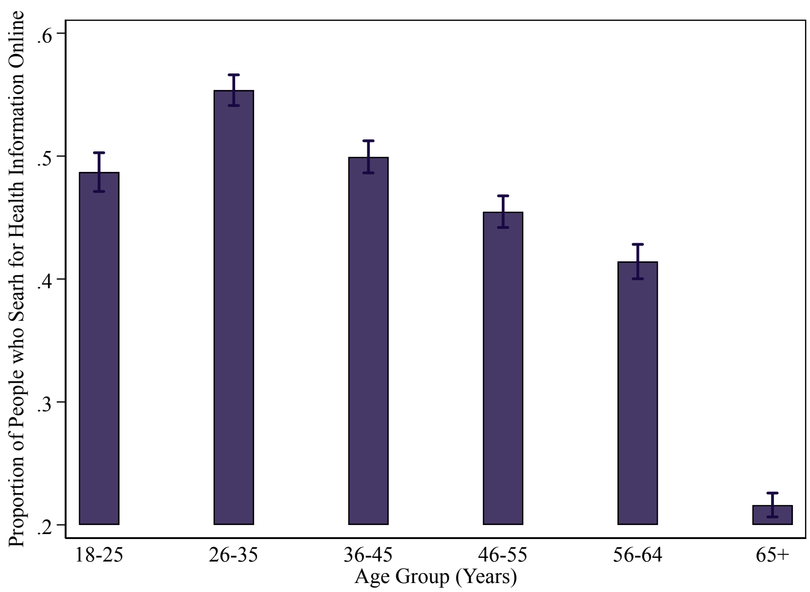

Figure 1.7. Internet Search by Age
Notes: This figure plots the proportion of 2011 NHIS respondents who reported having looked up health information on the internet in the past twelve months.

1.7. Data – Switchers

I now turn to the switcher analysis to understand the downstream outcomes of switching to higher quality, lower cost physicians. This analysis requires a panel of patients who are each matched to a single physician. To create this panel, I use claims data from patients who were insured between 2015 and 2021. Because effects are limited to primary care physicians, I limit my analysis to specialists in family medicine and internal medicine who were subjected to Insurer X's disclosure program.[14] Each patient is matched to their modal primary care physician: the physician they saw for the most visits in each year. Approximately 98% of patients had two or more modal primary care physicians. I keep these in the analysis dataset and drop any patients who did not visit any physician over the full seven years included in the dataset.[15]

To create a more exogenous measure of physician switching, I keep only patients who switched modal physicians in the same year that their original physician either exited the dataset (retired or became out-of-network for the insurer) or moved far enough away that the first two digits of their zip code changed. These patients were likely forced to switch physicians as opposed to having made the decision to switch based on some characteristic related to their outcomes. The final dataset includes 41,279 observations for 5,897 patients over seven years.[16]

This analysis explores the impacts of quality and cost on both quality and cost–related measures of healthcare utilization. I first study measures of spending, exploring the impact of physician quality and cost on the total amount paid by both the patient and the insurer for care provided by the modal primary care physician. I also explore the total amount paid by both the insurer and the patient over all physicians the patient saw in a year to determine whether there are any spillovers in cost. To determine whether patients are

[14]Program participation is largely determined based on specialty and geographic area. If an internal medicine physician primarily practices in a subspecialty that is not covered by the program (say, palliative medicine) they are not eligible for evaluation in the program

[15]I examine robustness of results using a panel that is balanced on the year relative to the switch rather than the calendar year.

[16]This dataset includes patients who interacted with the health system every year (though continuous enrollment is not specifically required), whose modal primary care physician is eligible for the disclosure program, who only have one primary care physician and see an eligible primary care physician each year. Appendix Table A.5 displays the total number of patients for each of these subsets.

better or worse off financially after having switched physicians, I also examine the total out-of-pocket payments made by the patient for their modal primary care physician and overall physicians.

To understand whether impacts on spending based on quality and cost are driven by volume of services or the price of services, I explore impacts on four outcomes: the number of services the patient received that year (measured by the number of unique service lines on the medical claim), the number of visits the patient had with their modal physician, the average number of services per visit (dividing the number of services by the number of visits), and the average "price per visit," which is the amount paid by the insurer and the patient divided by the number of visits.

Next, I examine emergency department utilization. I follow Alexander et al. (2019) by separating emergency room visits into three categories: visits for true emergencies that are preventable, to some extent, by appropriate primary care (" preventable emergencies"), visits for non-emergency care (" unnecessary visits"), and visits for true emergencies that are not primary-care preventable (called "placebo visits," since these should not change with higher quality or lower cost primary care).

Preventable emergencies are measured as any emergency room visit[17] for asthma, diabetes, influenza, heart attack, angina, or stroke. Asthma attacks, diabetes complications, and heart attacks can all be prevented to an extent by appropriate prescriptions and training from the physician on how to manage the condition, and influenza emergency visits may be prevented by an annual flu vaccine (Armstrong, 2014; Li et al., 2005; Hermida

[17]Emergency room visits include any claim with either an emergency room place of service code (23), an emergency room revenue center code (450–458), or emergency room CPT codes (G0380–G0384). These capture emergency room visits regardless of whether the patient was subsequently admitted to the hospital.

et al., 2011). Of course, much of the management falls to the patient, but part of being a higher quality physician is being able to teach patients how to appropriately manage their chronic conditions.

Although preventable emergencies are generally regarded as negative outcomes, it is possible that higher quality physicians could increase preventable emergencies by changing the threshold at which a patient decides an emergency visit is necessary. For example, a high-quality physician might communicate clear guidelines about when an emergency room visit is appropriate, while a lower-quality physician may not. Under the care of a higher quality physician, this could result in a patient going to the emergency room more frequently, if their health issues were always above the threshold for seeking emergency care, but the patient did not know that until after having seen the higher quality physician.

I measure unnecessary visits as emergency room visits for primary-care treatable conditions: urinary tract infections, conjunctivitis, upper respiratory tract infections, sore throat, and ear infections. These are visits that would most appropriately be treated in a primary care or office setting rather than an emergency department. Physicians who are of higher quality and who are lower cost should decrease these visits, redirecting care to the appropriate setting.

To measure placebo visits, I count the number of visits due to childbirth, poisonings, and fractures. These conditions need to be treated in an emergency room setting but are not outcomes a physician has control over. No increase or decrease is expected with a switch to a higher-quality or lower cost physician.

Table 1.6 shows summary statistics over all outcomes for physicians broken into three categories of patients: "stayers," those who continued to see the same primary care physician over the full panel, "switchers," who switched physicians at least once, and "induced switchers" whose physician switch coincided with the year their former physician moved or exited the dataset. The first nine rows in Table 1.6 display the means and standard deviations of spending and emergency room utilization variables, and the following rows display differences in physicians' average quality and cost scores as well as the average change in these scores for patients who switch physicians. Across all outcomes and quality and cost variables, stayers, switchers, and induced switchers look similar.

Table 1.6. Summary Statistics

	Stayers	Switchers	Induced
Total Paid	713.804	681.288	680.496
	(779.172)	(1136.224)	(747.094)
Services per Visit	2.272	2.208	2.202
	(1.788)	(1.77)	(1.74)
Services per Visit	2.272	2.208	2.202
	(1.788)	(1.77)	(1.74)
Price per service	131.242	145.174	142.588
	(90.206)	(109.826)	(103.032)
Total Paid All	5117.684	5654.046	5553.242
	(14130.74)	(15414.66)	(15922.63)
Non-modal Doctor Paid	4403.878	4972.758	4872.746
	(14023.25)	(15267.15)	(15823.38)
Preventable ED	.022	.022	.026
	(.512)	(.456)	(.514)
Non-Emergency ED Visits	.006	.008	.004
	(.166)	(.156)	(.104)
Placebo ED Visits	.002	.002	.002
	(.256)	(.244)	(.164)
ΔQ	0	-.816	-.466
	(0)	(17.396)	(16.09)
ΔC	0	-.092	-.254
	(0)	(2.824)	(2.684)
ΔQC	0	.13	-.86
	(0)	(45.604)	(40.248)
Observations	1302399	613802	41279

Notes This table displays summary statistics for the dataset used in downstream outcomes analysis. The average and standard deviation are displayed for each outcome variable and for right hand side variables. Patients are broken into three groups. Stayers are those patients who see the same primary care physician over the duration of the panel. Switchers are those who switch physicians, and Induced patients are those whose switch was induced by their former PCP either leaving or moving.

1.8. Empirical Model – Switchers

I exploit patient switches between physicians to determine the impact of quality and cost on utilization outcomes. Adapting the model from Finkelstein et al. (2016) to my setting, I define the patient's utility as a function of their utilization conditional on health status h_i and a preference parameter η_i. There is an appropriate level of utilization (y) given the patient's health status that is adjusted by the preference parameter. A patient with a higher level of η_i prefers more utilization; perhaps they use the emergency department more frequently or prefer physicians to prescribe more invasive or expensive procedures.

$$u(y \mid h, \eta) = -\frac{1}{2}(y - h_i t)^2 + \eta_i y \tag{1.4}$$

The physician maximizes her own utility, $\tilde{u}$, which includes patient utility adjusted by the physician's own quality and cost parameters. Both quality, λ_j^q, and cost, λ_j^c, impact how much care physicians give to their patients and may impact care in different ways. For example, a higher quality physician may run more tests, leading to higher utilization, whereas a more lower cost physician may choose less expensive tests. By assumption, quality and cost are separable.[18] The physician also faces time-varying costs of providing care, PC_{jt}.

[18]The separability assumption can be relaxed by modeling the utility as $u = -1/2(y - h_{it})^2 + \eta_i y + \gamma_j^c y + \gamma_j^q (1 - lambda_j^c) y - PC_{jt}$. The interaction of λ_j^q with $(1 - \lambda_j^c)$ can be interpreted as higher-quality physicians leading to higher utilization (perhaps through increased testing), while higher-quality physicians are also less biased away from the "optimal" level of utilization (that which maximizes $-1/2(y - h_{it})^2$) by their own cost. In other words, a lower-quality, lower cost physician might skimp on all tests, whereas a higher-quality, lower cost physician would run only the important tests. Results in Appendix Table 5 imply no such complementarities exist in the data, so the separability assumption is likely valid.

$$\tilde{u} = -\frac{1}{2}(y - h_i t)^2 + \eta_i y + \lambda_j^c y + \lambda_j^q y - PC_j t. \tag{1.5}$$

The physician chooses the level of provision to maximize her utility function. The physician's optimal choice of utilization is a function of patient-specific parameters, physician cost and quality, and a time-varying component. I assume that PC_{jt} is linear in utilization and additively separable in j and t, and that the physician component of PC_{jt} is captured by λ_j^c. This assumption implies that the physician's cost parameter reflects both the physician's own propensity to provide lower cost care and how costly it is for that physician to provide care. The separability assumption allows PC_{jt} to be displayed as the linear combination of λ_j^c and a time fixed effect τ_t. I next assume that the level of utilization that maximizes $-1/2(y - h_{it})^2$ is comprised of a patient fixed effect, α_i, and a set of observable time-varying controls, x_{it}, which include relative switch-time indicators. These indicators capture variation in optimal utilization based on the amount of time until a switch, allowing for the decision to switch physicians to be correlated with changes in health status over time. Under these assumptions, first-order conditions from a maximization of the utility in Equation 1.5 imply that utilization y_{ijt} for a patient i who saw physician j in year t is:

$$y_{ijt} = \hat{\alpha}_i + \tau_t + \lambda_j^c + \lambda_j^q + x_{it}\beta + \epsilon_{ijt}. \tag{1.6}$$

I then rewrite Equation 1.6 as a function of pre- versus post-switch timing in Equation (7) as follows:

$$y_{it} = \hat{\alpha}i + \tau_t + \lambda o^c + \lambda_o^q + 1\{\text{Post}\}(\theta^Q \Delta Q + \theta^C \Delta C) + x_{it}\beta + \epsilon_{it}. \tag{1.7}$$

where $\Delta_Q = Q_d - Q_o$ and $\Delta_C = C_d - C_o$, the difference in physician scores for quality and cost between the "destination physician," (d), that the patient switched to and the "origin physician," (o), that the patient switched from (cost scores are negated so that an increase in Δ_C refers to a decrease in physician costliness). Note that in Equation 1.7, the physician scores (Q, C) are not equivalent to the physicians' costliness and quality parameters λ_Q, λ_C. These scores instead include both physician-specific and patient-specific components of measurable costliness and quality. The parameters θ^Q and θ^C in Equation 1.7 therefore represent the ratio of the difference in quality and costliness to the difference in measurable costliness or quality scores: $\theta_c = \frac{\lambda_d^C - \lambda_o^C}{C_d - C_o}$. These parameters can then be interpreted as the increase in utilization that is attributable to physicians per a one-unit increase in costliness or quality scores. These need not lie between zero and one because the units of the numerator and denominator are different.

Finally, note that the characteristics of the origin physicians, λ_o^C and λ_o^Q, are captured empirically by the patient fixed effect, so that the regression I run is displayed in Equation 1.8,

$$y_{it} = \alpha_i + \tau_t + 1_{\text{Post}}(\theta^Q \Delta Q + \theta^C \Delta C) + x_{it}\beta + \epsilon_{it}. \tag{1.8}$$

where $\alpha_i = \hat{\alpha}i + \lambda o^C + \lambda_o^Q$.

The model implicitly assumes that outcomes depend only on the current physician and not the history of physicians that the patient has seen. For most patients, this assumption is satisfied because the utilization outcomes studied are quick to come to fruition. For example, corticosteroids used for the treatment of asthma take only four to six weeks to improve breathing.[19] If the assumption is not met, then θ^Q and θ^C simply measure deviations from the lasting impacts of the prior physicians, which are absorbed by the patient fixed effects. The model also implicitly assumes homogeneity in treatment effects: that θs do not vary with patient characteristics. This assumption will be relaxed in various specifications that explore impacts within older and younger patients.

Recent literature has shown that heterogeneity over treatment timing, where earlier-treated groups respond differently than later-treated groups, can also bias results. In Appendix Section A.2, I explore robustness to this concern using novel estimators that relax the assumption of no heterogeneity over treatment timing.

To interpret θs as causal parameters, the post-switch indicator as well as Δ_Q, and Δ_C must be exogenous. This means that there cannot be any unobservable time-varying characteristics that correlate with both the outcome and either the post-switch indicator or the change in quality or costliness of the physician. There is evidence in this study that different types of patients are more or less likely to choose high-quality, lower cost physicians; however, these patient types are generally fixed over time. Therefore, the correlation is captured by the patient fixed effect. To increase the plausibility of exogeneity, I

[19]https://www.mayoclinic.org/drugs-supplements/corticosteroid-inhalation-route/proper-use/drg-20070533

focus on the subset of patients who switch physicians because their previous physician either moved or exited the dataset. I call these "induced switches." These induced switches were forced, ruling out patient selection in the decision to change physicians.

While the assumption of parallel pre-trends is not necessary or sufficient for the interpretation of thetas as causal parameters (Hull, 2018), it is still helpful to visualize trends in spending for patients switching to higher- versus lower-quality and lower cost physicians. If trends were markedly different leading up to the switch year for those who switched to higher-quality physicians, it would be more difficult to justify the assumption that there is no selection on time-varying unobservable characteristics into higher-rated physicians. To visualize pre-trends, I regress the following, estimating separate parameters for each year relative to each patient's first switch.

$$y_{it} = \alpha_i + \tau_t + \sum_{k=-6}^{5} 1\{\text{Year} = k\}(\theta_k^Q \Delta Q + \theta_k^C \Delta C) + x_{it}\beta + \epsilon_{it}. \tag{1.9}$$

1.9. Impacts on Downstream Outcomes

In this section, I explore the results of the switcher regressions in Equations 1.8 and 1.9, both on average and broken down by patient age.

1.9.1. Average Effects

I estimate Equations 1.8 and 1.9 for patients who saw primary care physicians between 2015 and 2021 on measures of utilization. First, I explore total spending, analyzing the amount paid by both the insurer and the patient each year. Figure 1.8 displays the results of estimating Equation 1.9 on total spending. Panel A reports the coefficients on a one-point increase in lower cost, while Panel B reports the coefficients on a one-point increase in quality. In both panels, pre-trends are relatively flat, with individuals who switch to lower cost physicians trending similarly, and likewise for higher quality. After the patient switches physicians, spending decreases relatively for every single-point decrease in cost scores and increases relatively for every additional quality point. Figure 1.8 includes data on all patients who saw a primary care physician who was part of the disclosure program during each calendar year; however, this mechanically means that the panel is unbalanced on time relative to the switch. To overcome this issue, Appendix Figure A.9 displays the impacts using a dataset that is instead balanced on relative time: the years leading up to and after a switch. Balancing on relative year removes numerous observations, so effects are underpowered; however, they are similar in magnitude, so it is unlikely that panel imbalance on relative year biases effects.

Table 1.7 shows the results of the regression in Equation 1.8 on a variety of spending outcomes. Column 1 reports the impacts of a single-point increase in quality and for a single-point decrease in cost. A single-point increase in quality increases payments by \$1.65, while a single-point decrease in cost decreases payments by \$11.29. Appendix Table A.6 shows the results of these regressions including an interaction term between cost and quality scores, and there is no significant impact from the interaction, adding credence to the assumption that cost and quality are separable.

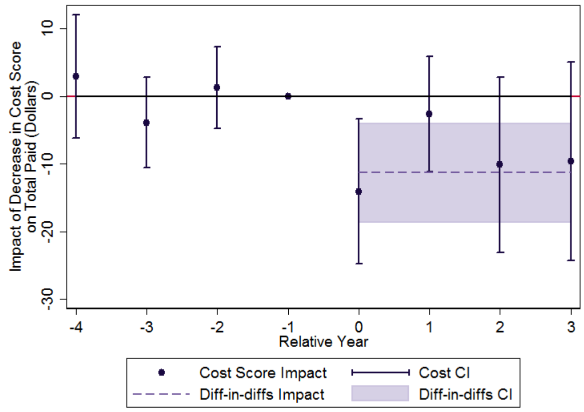

(a) Effect of Cost

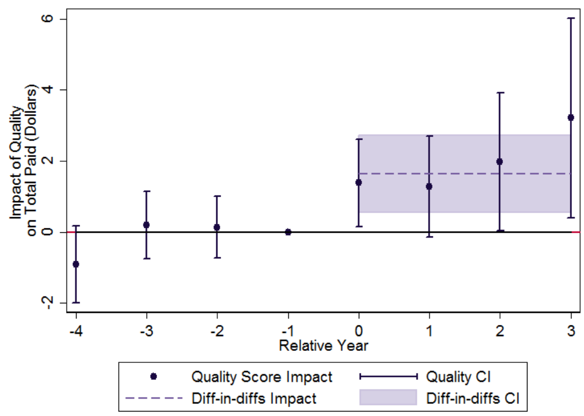

(b) Effect of Quality

Figure 1.8. Effect of quality and cost scores on total amount paid
Notes: This figure displays the impacts of switching to a single point lower-cost physician (panel A) or higher quality (panel B) physician on the total amount paid: the sum of patient- and insurer spending.

	(1) Total Paid	(2) Services	(3) Services per Visit	(4) Price per service	(5) Total Paid All	(6) Non-modal Doctor Paid	(7) Patient Paid
$\Delta Q\times$ Post	1.646***	0.0178***	0.00537***	-0.120	16.36	14.71	0.338
	(0.553)	(0.00676)	(0.00157)	(0.0761)	(11.23)	(11.22)	(0.238)
$\Delta C\times$ Post	-11.29***	-0.0104	0.0103	-1.641***	-107.3**	-96.03**	-2.992*
	(3.707)	(0.0415)	(0.00825)	(0.522)	(45.36)	(45.06)	(1.761)
R-Squared	0.415	0.510	0.508	0.547	0.452	0.450	0.453
Outcome Mean	680.5	6.34	2.2	142.58	5553.24	4872.74	186.38
Average Impact	-16.98	.02	.04	-3.24	-159.72	-142.74	-4.7
Observations	41279	41279	41279	41279	41279	41279	41279

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.7. Effects of Quality and Cost on Spending Outcomes

Notes: This table shows the impacts of quality and cost scores on spending outcomes. Cost scores are negated so that estimates represent the effects of a single-point decrease in cost scores. Column 1 reports the impacts on total spending: the amount paid for visits with the modal primary care physician by both the patient and the insurer. Column 2 reports impacts on the number of service received by the modal primary care physician. The outcome in column 3 is the average number of services received per visit, and column 4 reports the average price per service: total amount paid divided by total number of services. Columns 5 and 6 explore possible spillovers in spending on other providers. The outcome for the regression in column 5 is spending over all providers, not just the modal primary care provider. Column 6 narrows down to spending over other providers, not including the modal primary care provider. Column 7 displays impacts on the portion the patient paid out of pocket.

To put these impacts in perspective, when the average patient switches from "lower quality" to a "higher quality, lower cost" physician, the quality score increases by 2.18, and the cost score decreases by 1.82. Multiplying these by the impacts of single-point increases predicts that for the average switch from a "lower quality" to a "higher quality, lower cost" physician, a patient saves \$16.98.

Table 1.7 also examines mechanisms by which spending may increase or decrease, particularly by breaking down total spending into quantity (number of services received) versus price (amount spent per service received). Columns 2 and 3 in Table 1.7 display impacts of switching to higher- quality and lower cost physicians on the number of services and services per visit, while column 4 shows the impacts on the average price per service: the total amount paid divided by the number of services.

Interestingly, the mechanisms for impacts on spending are different for quality and cost. Spending increases that are associated with switches to higher-quality physicians are due to increases in both the number of services in total and on a per-visit basis, with single-point increases in quality leading to small but significant increases in the number of services the patient received, but if anything, slightly less expensive services being rendered. On the other hand, physicians with lower cost scores seem to achieve these lower scores by charging lower prices or choosing less expensive treatments, where single-point decreases in cost scores lead to \$1.64 less being spent per service rendered.

Columns 5 and 6 of Table 1.7 display impacts on the total amount paid over all physicians an individual saw each year: column 5 shows spending on all physicians, and column 6 shows the amount paid for services received by physicians other than their modal PCP. The results in these columns show that when a patient switches to a lower cost physician, their overall spending across all physicians decreases. This is not just because the PCP's contribution to total spending is high but because spending associated with non-modal physicians also decreases. This could be because physicians who are lower cost refer to specialists who are lower cost. However, it could also be that the lower cost physicians suggest less expensive procedures, and other physicians who actually perform those procedures are noted on the claim rather than the original physician.

The final column of Table 1.7 displays the impacts on patients' out-of-pocket spending, which is a subset of total spending. The results in this column show that switching to a lower cost physician results in a marginally significant decline in out-of-pocket spending. The proportional impact is the same as the impact of lower cost on total spending: the $11 decrease on a base of $680 in total spending is almost identical proportionally to a $3 decrease on a base of $186. Thus, these results do not imply that the insurer receives all the gains to lower cost; more likely, the magnitudes of the effects on patients' out-of-pocket spending are too small to be detected by the sample size.

I next explore outcomes on emergency room utilization. In Table 1.8, I report the results of estimating the regression in Equation 1.8 on three measures of emergency room utilization: preventable emergencies, non-emergencies, and placebos. Column 1 shows the impacts on preventable emergency room visits, which are true emergencies that are in part preventable with high-quality primary care. Column 2 shows impacts on unnecessary emergency room visits, which are not emergencies and are more appropriately treated in an office setting, and Column 3 shows impacts on placebo emergencies, which are true emergencies that cannot be decreased or increased through improved primary care.

	(1) Preventable ED	(2) Unnecessary ED	(3) Placebo ED
$\Delta Q\times$ Post	0.0000378	0.0000645	-0.00000824
	(0.000267)	(0.0000878)	(0.0000858)
$\Delta C\times$ Post	0.0000397	0.000730	0.000475
	(0.00146)	(0.000578)	(0.000519)
R-Squared	0.361	0.194	0.160
Outcome Mean	.04646	.0126	.00982
Average Impact	.00016	.00146	.00084
Observations	41279	41279	41279

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.8. Effects of Quality and Cost on Emergency Department Utilization

Notes: This table shows the impacts of quality and cost scores on emergency room utilization outcomes. Cost scores are negated so that estimates represent the effects of a single-point decrease in cost scores. Column 1 displays the impacts on preventable emergency department (ED) visits, which are conditions that are true emergencies but which appropriate primary care could have in part prevented. Column 2 explores impacts on non-emergency ED visits, which are non-emergencies that would more appropriately be treated in a primary care office setting. Column 3 reports impacts on "placebo" emergency department visits, which are true emergencies that are not preventable by higher quality or lower cost preventative care.

Table 1.8 shows no evidence of any statistically significant impacts of quality or cost on emergency room utilization. The 95% confidence intervals on the effect of a one-point increase in quality range from a decrease of 0.0005 emergency room visits to an increase of 0.0005, and the effect of a one-point decrease in cost range from a decrease of 0.003 in emergency room visits to an increase of 0.003. These confidence intervals are fairly precise, including a 1.2% increase and 1.0% decline in preventable emergencies for a physician who has one-point higher quality and an 6.2% increase and an 6.0% decrease in preventable emergencies for physicians who are one-point lower on cost. Null effects on preventable emergencies can also be seen graphically in Appendix Figure A.14.

One concern about using preventable emergencies as the main health outcome is that these rare outcomes may be difficult to move. However, other interventions have been shown to significantly impact these outcomes. Alexander et al. (2019) find that retail clinics can significantly decrease both preventable and unnecessary emergency room visits, and Miller (2012) finds that having health insurance leads to fewer unnecessary emergency room visits. To improve the statistical power of results, I also explore impacts when including endogenous switches (see Appendix Table A.9). My findings show that these impacts are similar in magnitude, but the effects are more precise, with confidence intervals on quality spanning a 0.3% decrease to a 0.2% increase in emergency room visits and confidence intervals on cost spanning a 1.8% decrease and a 1.5% increase.

1.9.2. Heterogeneous Effects

I next separate patients by age group, calculating the impacts of changes in quality and cost on outcomes separately within each group. Table 1.9 displays the impacts for six 10-year age bins for patients aged 18 and older. Column 1 shows impacts for the youngest patients, aged 18-25, and different age groups up to Column 6, which shows impacts for those who are age 65 and older.

Columns 1 and 2 show no significant impacts of quality or cost on total spending for the younger patients, while columns 3 and 4 show both significant increases in spending due to increased quality and significant declines in spending due to decreases in cost, respectively. The impacts then decline again for patients who are aged 56–64, and while patients aged 65 and older do show large point estimates, none are statistically distinguishable from zero. Multiplying the impact of cost and quality by the changes in quality and cost scores for the average (within-age-group) switch from a "lower quality" to a "higher quality, lower cost" physician means that patients who are 46–55 years old stand to save nearly

	(1) Total Paid	(2) Total Paid	(3) Total Paid	(4) Total Paid	(5) Total Paid	(6) Total Paid
$\Delta Q\times$ Post	5.493	2.465	3.920**	2.214**	0.590	0.618
	(4.086)	(3.407)	(1.584)	(0.955)	(0.667)	(3.262)
$\Delta C\times$ Post	1.549	5.273	-16.44**	-19.66**	-4.342	-13.03
	(17.13)	(11.36)	(8.336)	(7.635)	(5.639)	(10.78)
Age Range	18-25	26-35	36-45	46-55	56-64	65 and older
Average Impact	.28	13.2	-16.14	-29.64	-7.06	-26.58
Standard Error	(26.6)	(19.2)	(11.66)	(13.66)	(11.02)	(25.18)
Average Outcome	635.68	610.7	660.96	690.9	727.6	586.06
R-Squared	0.432	0.437	0.435	0.366	0.446	0.420
Observations	658	1463	4648	11445	16121	6944

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.9. Effects on Total Spending By Age Group

Notes: This table shows the impacts of quality and cost scores on total spending, broken into different age groups. The impacts for the youngest adult patients, aged 18-25 are shown in column 1, and age increases over the columns with the oldest patients, aged 65 and above, shown in column 6.

$30 in the year of the switch, whereas the youngest patients would see a not statistically significant spending increase of $0.28.

The number of patients in the dataset varies with age, with only 658 observations in the youngest age group, and up to approximately 16,000 in the older age groups. One may wonder whether results for the younger patients are too underpowered to detect significant changes in spending. One way to examine this hypothesis is to include all switches, including the switches that are not induced by physician moves or exits. The assumption that switches are exogenous is less plausible for this larger sample; however, there are many more switches to learn from. Appendix Table A.7 displays the results from running the switcher regression on the full set of physician-patient switches. Impacts of cost for all age groups are larger, so the inclusion of endogenous switches may bias results away from zero. However, the pattern of effects is the same, with the largest impacts for the middle-aged patients who are aged 36–55, and smaller impacts for other patients.

The general pattern of results remains when analysis is done on the subset of the data that is balanced on year relative to the switch rather than calendar year. The subset contains patients who exist in the dataset from two years before the switch until two years post-switch. The requirement of having two years post-switch included in the data means that only switches from 2017–2019 can be included (i.e., 2020 and 2021 are the two post-switch years; 2015 and 2016 are pre-switch years). This decreases the size of the dataset considerably but removes any concern that longer-run impacts are driven by changes in the sample rather than true dynamic effects. Appendix Table A.8 displays impacts over age groups on this subsample. The general pattern of results remains, with the patients aged 36–45 experiencing the largest benefits to switching to lower cost physicians. However, the impacts of physician quality are less robust.

In a final robustness check, Appendix Figure A.10 shows the dynamic effects within age groups. The purpose of this check is to ensure that within age groups, patients who switch physicians earlier are otherwise trending similarly to patients who switch later. This check rules out the possibility that different effects are driven by different trends within age groups.

Nonlinearity over patient age may also be present for health outcomes, even though no effects on preventable emergencies were found on average. Table 1.10 shows the impacts by age group on preventable emergency room visits. There are no statistically significant impacts of quality on preventable emergency room visits, regardless of a patient's age. There are two age groups in which statistically significant impacts on cost are detected. Patients aged 46–55 experience increases in preventable emergencies when they switch to lower cost physicians, and patients aged 56–64 experience marginally significant fewer preventable emergencies.

	(1) Preventable ED	(2) Preventable ED	(3) Preventable ED	(4) Preventable ED	(5) Preventable ED	(6) Preventable ED
$\Delta Q\times$ Post	-0.00000117	-0.000850	0.00157	-0.000549	0.000110	0.00125
	(0.000401)	(0.00107)	(0.00104)	(0.000499)	(0.000318)	(0.000883)
$\Delta C\times$ Post	0.00512	-0.0123	0.00132	0.00764***	-0.00515***	-0.00104
	(0.00583)	(0.0129)	(0.00243)	(0.00292)	(0.00192)	(0.00463)
Age Range	18-25	26-35	36-45	46-55	56-64	65 and older
Outcome Average	.08	.04	.02	.02	.02	.02
Average Change	0	-.02	.004	.012	-.01	0
Standard Error	(0)	(.02)	(.004)	(.006)	(.004)	(.01)
R-Squared	0.243	0.332	0.355	0.384	0.319	0.395
Observations	658	1463	4648	11445	16121	6944

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.10. Effects on Preventable Emergencies By Age Group

Notes: This table shows the impacts of quality and cost scores on preventable emergency department visits, broken into different age groups. The impacts for the youngest adult patients, aged 18-25 are shown in column 1, and age increases over the columns with the oldest patients, aged 65 and above, shown in column 6.

I again examine robustness of these results to both inclusion of endogenous switches in Appendix Table A.9, and make a subset of a relative-time balanced panel in Appendix Table A.10. Appendix Table A.9 shows that when all switches are analyzed, neither the marginally significant decline in preventable emergencies for patients aged 56–64 nor the significant increase for patients aged 46–55 remain. Additionally, there is a marginally significant increase in preventable emergencies for the youngest patients. Appendix Table A.10 shows that the decline in emergency room visits for ages 56–64 is again robust; however, for ages 46–55, the significant increase in preventable emergencies is no longer there, while a new significant decline in preventable emergencies as a function of quality shows up for the same age group (for a one-point increase in the quality score). Together, these exercises suggest that the decline in preventable emergency room visits for patients aged 56–64 is robust, but the increase in preventable emergency room visits from switching to lower cost physicians for patients aged 46–55 is possibly driven by changes in dataset composition (i.e., denominator) over years post-switch. Overall, the analysis points to decreases in spending being larger for middle-aged patients with marginally significant decreases on preventable emergency room visits for the 56–64 age group.

1.9.3. Mechanisms

As we have seen, the effects of underlying physician quality and cost are concentrated on middle-aged patients; there are a few reasons why this may make sense. First, middle-aged patients are likely to be newly diagnosed with a chronic illness. Appendix Figure A.11 shows that the likelihood of having a current diagnosis that did not exist in the previous year for any chronic illness has an inverted U-shape over age, peaking between ages 40 and 60. Table 1.11 explores whether the onset of chronic illness explains the fact that middle-aged patients are more heavily impacted by the quality and cost of their physicians, separately estimating impacts over the age distribution for patients who became newly chronically ill during the sample period. Column 1 of Table 1.11 displays impacts on the amount paid for younger patients, aged 18–35, who have no change in their health status, whereas column 2 displays results for those who became chronically ill at some point during the sample period. Columns 3 and 4 display these results for middle-aged patients, while columns 5 and 6 display the same for the oldest patients.

	(1) Total Paid	(2) Total Paid	(3) Total Paid	(4) Total Paid	(5) Total Paid	(6) Total Paid
$\Delta Q\times$ Post	9.601**	-0.918	0.763	4.280***	0.553	0.465
	(3.988)	(3.360)	(0.956)	(1.057)	(1.751)	(0.874)
$\Delta C\times$ Post	17.78	-4.011	-20.44	-18.13***	-8.492	-6.807
	(11.63)	(13.43)	(14.50)	(5.184)	(9.282)	(5.872)
Age Range	18-35	18-35	36-55	36-55	56 and older	56 and older
Health Status	No Change	Newly Chronically Ill	No Change	Newly Chronically Ill	No Change	Newly Chronically Ill
R-Squared	0.459	0.413	0.450	0.357	0.402	0.480
Observations	1078	1043	5509	10584	8596	14469

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.11. Effects on Amount Paid By Age Group and Chronic Illness Status

Notes: This table shows the impacts of quality and cost scores on preventable emergency department visits, broken into different age and chronic illness groups. The impacts for the youngest adult patients are in columns 1 and 2, and age increases over the columns, with the oldest patients shown in columns 5 and 6. The columns also break down patients by chronic illness status, with columns 1, 3, and 5 displaying impacts for patients without any change in chronic illness status over the time period studied, and columns 2, 4, and 6 displaying impacts for those with a new or recent chronic illness diagnosis.

While the onset of chronic illness is correlated with age in the same way treatment effects are, it does not explain the pattern of results over patient age. If the onset of chronic illness did explain results, then one would expect to see significant effects across the age distribution for the newly chronically ill. This is not the case. There is some evidence that the impact of cost for middle-aged patients is driven by the newly chronically ill; however, the point estimate for healthy middle-aged patients is just as large but noisier. There are also some differences noted in the impact of quality on spending that differ by health status, but these patterns do not imply that health status explains the pattern of results over patient age.

There is no evidence that the onset of chronic illness explains the patterns over patient age. However, there are many other characteristics that vary with age and could instead explain the results, many of which cannot be measured without additional data. For example, younger patients in this dataset may be more diligent about their health (they visit their physician's office annually despite being quite healthy), whereas older patients may be less diligent and therefore more reliant on their physician to make good treatment choices. Age could also vary with factors such as income or socioeconomic status (SES) to the extent that selection into private insurance is impacted by these factors differently by age. Diligence and SES could therefore play a role in explaining the differential impacts by age.

1.10. Limitations

This paper uses two main identification strategies: regression discontinuity (RD) and a two-way fixed effects "switcher" model. Interpreting effects from both models requires several assumptions, and certain limitations must be recognized.

First, regression discontinuity estimates are local around the RD cutoff: for physicians further from the cutoff, the effects of the program may differ. Second, the analysis focuses on physicians who were assigned "higher quality, lower cost" status or "higher quality" status based on their scores, leaving out physicians who were assigned their status because they had a score that was valid for two years and wasn't updated in 2019. Some other physicians had missed the quality or cost cutoff, but upon a reconsideration request were granted "higher quality, lower cost" or "higher quality" status. For these physicians, their original status (pre-reconsideration request) is used in the analysis, so as not to bias results due to selection into making a reconsideration request. However, this means that effects may not generalize to those who achieved a different status based on a reconsideration request later on if that group of physicians is different in ways that impact the effect of having a higher status.

Note further that some individuals are granted "higher quality, lower cost" status because their group achieves cost, even if they themselves do not. These physicians are included in the analysis, which is why the first stage plot in Appendix Figure A.5 (b) shows that fewer than 100% of physicians passing the cutoff achieve "higher quality" status (since some are given "higher quality, low cost" status due to their group's achievements instead). It is also why Appendix Figure A.5 (c) shows that some physicians who did not meet the cost criteria are given the "higher quality, lower cost" status. This means that effects should be interpreted as "intent-to-treat" effects: the effect of passing the cutoff, rather than as "true" causal effects of being assigned a higher status. This limitation can be partially overcome by using instrumental variables (IV) analysis as discussed in Section 1.5; however, it is still possible that these IV analyses are under-powered to detect the effect of comparing "lower quality" to "higher quality" and "higher quality" to "higher quality, lower cost."

Next, it is important to note that as with many programs put in place by private insurers, the program iterates year over year. This paper focuses on the 2019 version of the program, and changes may have been made since then which would impact the accuracy of extrapolation of results to future iterations of the program.

The two-way fixed effects "switcher" results should also be interpreted knowing some important limitations. First, while the effects from the "switcher" model can be considered causal effects under the assumptions discussed in Section 1.8, the comparison of effects between models (and for example, between age groups) should be considered descriptive rather than causal. That is to say, there are many reasons why older patients may have higher potential cost savings than younger patients. Section 1.9 explores the possibility that the differences in effects are driven by the onset of chronic illness, but there are other possible explanations as well, such as that younger patients may choose different types of plans that impact how much influence a physician can have over their spending. Further, since younger patients have lower healthcare utilization in general than older patients, effects on younger patients may be underpowered (there are many fewer observations for younger patients). Finally, the absence of effects on emergency department visits could be driven by a lack of statistical power, and this possibility is explored at the end of Section

1.9.1, where effects on all switchers (not just those induced to switch by a physician move) are analyzed to improve statistical power. Of course, there is a tradeoff in this analysis: while including non-induced switches improves the power of the analysis, it may also bias effects by allowing for patient selection into switching.

1.11. Discussion and Conclusion

This study aimed to answer two research questions: (1) which patients respond to disclosure programs, and (2) which patients benefit most from responding to disclosure programs. My findings show that younger patients respond more to the signal of higher quality and lower cost than do older patients, so the average age of new patients seen by "higher quality, lower cost" physicians is significantly lower than the average age of patients seen by "lower quality" physicians. On the other hand, it is the middle-aged patients, not the youngest patients, who experience the largest cost savings from seeing higher-quality, lower cost physicians.

Figure 1.9 compares the estimates of the two sections of the paper, first plotting average patient response to physician status by age, showing that patients who are aged 18–35 respond the most to disclosure programs. "Higher quality, lower cost" physicians attract relatively more new patients from these age groups than from other age groups. In the lighter colored bars, I plot the average savings from switching from a lower-quality to a "higher quality, lower cost" physician separately by age group and find that effects are largest for the middle-aged patients who are 36–55 years old. Clearly, the patients who benefit most by switching to "higher quality, lower cost" physicians are not the same patients who respond to physician status.

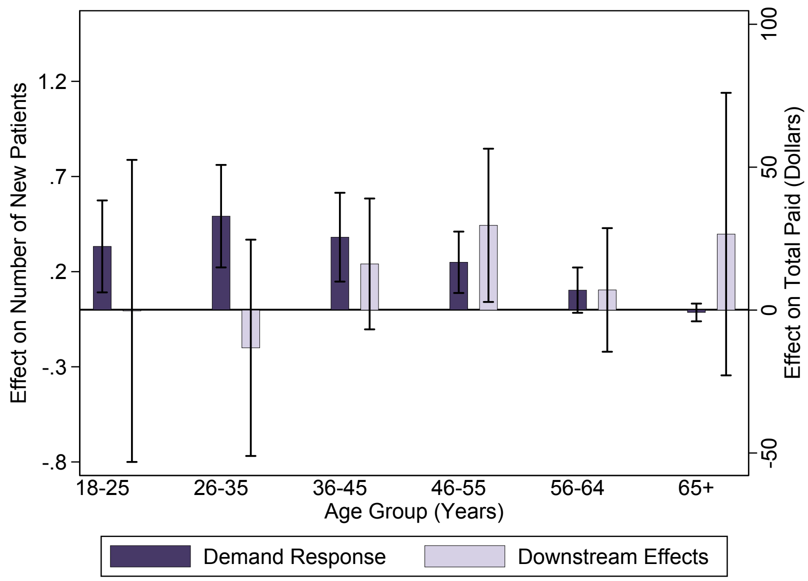

Figure 1.9. Comparison of Effects by Age Group
Notes: This figure plots, in dark purple, effects of the "higher-quality, lower-cost" status on the number of new patients in each age group. In light purple, the average savings from switching from a "lower-quality" to a "higher quality, lower cost" doctor is displayed separately for each age group. The patients who respond most strongly to the disclosure policy are those who save the least in total spending from switching.

Comparing results from the regression discontinuity design to those from the two-way fixed effects switchers design requires several assumptions. First, to interpret the regression discontinuity effects as the effect of being classified as "higher quality, lower cost" regardless of underlying quality level, one must assume that the treatment effect does not depend on the physician's quality or other characteristics that vary with the quality score. One may wonder whether quality correlates with capacity such that higher-quality physicians may experience lower gains from being classified as "higher quality, lower cost" physicians because they are already close to capacity. Although the findings show that impacts do vary with physician capacity (see Appendix Figure A.2), capacity does not vary substantially over the quality distribution. Second, one must assume that the young patients who are influenced by the disclosure program to choose "higher-quality, lower cost" physicians are comparable to the young patients in the switcher analysis panel (and likewise for all other age groups). The switcher panel includes patients who interact with the health system every year, so they are possibly sicker than young patients or are more diligent about their health (consistently going for annual checkups). On the other hand, inclusion in the response to physician status dataset does not require patients' repeated use of the health system, but does require them to have had an appointment in the three months after 2019 scores were released. When these assumptions are met, one can compare impacts and find that young patients respond to the "higher quality, lower cost" status, whereas older middle-aged patients experience larger savings.

Savings can be maximized when patients who have the most to gain are paired with the higher-scoring physicians. In this light, it would be wise to consider a policy that would re-sort patients so that the middle-aged patients (those who have the most to gain) are paired with the lower cost physicians. I simulate such a policy below. While this exercise ignores several important details (such as what type of policy could lead to this sorting, and other inputs to total welfare such as physician-to-patient distance and the patient's preferences over other physician characteristics), it serves to benchmark how much money could be saved by a policy that can achieve this type of maximal-savings, patient-physician matching.

To simulate such a policy, I begin with the dataset of all patients who switched physicians between 2015 and 2017. Within geographic areas, I re-assign patients to physicians based on their age, so that the middle-aged patients are assigned to the lower cost physicians, the oldest patients are assigned to the next lower cost physicians, and then the youngest patients are assigned to the least lower cost physicians. Each physician is assigned to the same number of patients that they treated in the original dataset, so capacity constraints are built into the exercise. I predict the change in outcomes from switching to the maximal-savings physicians by multiplying the age-specific effects in Table 1.9 by the changes in quality and cost status between the original physician and the maximal-savings–matched physician. After determining the potential savings from matching patients to their maximal-savings physician, I also determine potential savings from allocating patients randomly to a given physician (within geographic areas).

I find that a policy that could match patients with their maximal savings physician could save 7.5% in annual spending over and above the savings from the matches patients endogenously made under the current policy. When compared to a policy that randomly allocates patients to physicians the maximal savings matches save about 7.7% (increasing savings by 0.2%). While matches are not created based on health outcomes, the maximal-savings policy would also decrease preventable emergencies by about 2.6% of preventable visits relative to random matching (recall that the estimates on preventable emergencies were not statistically significant, so these estimates should be interpreted with caution). These numbers should be interpreted as order-of-magnitude estimates of what could occur under a policy that more effectively matches patients to physicians. However, these numbers should not be viewed as welfare estimates, since this analysis explicitly ignores patient preferences over their physicians.

This research draws several conclusions. First, privately insured patients respond to information about primary care physicians, and younger patients respond more. When patients do switch physicians, switching to lower cost physicians decreases patients' spending without significantly impacting outcomes, as measured by preventable emergency room visits. Middle-aged patients are impacted the most by their physician's quality and cost, while the youngest patients are more heavily influenced by the disclosure program and seek higher-quality, lower cost physicians. There is room for policy improvement through policies that focus information or financial incentives on middle-aged patients.

CHAPTER 2

Housing Affordability and Domestic Violence: The Case of San Francisco's Rent Control Policies

(with Eilidh Geddes)

2.1. Introduction

Domestic violence makes up 21% of all violent victimizations in the United States (Truman and Morgan, 2014). Four women are killed, and 14,000 are battered by their partners daily (Aizer, 2010). The prevalence and severity of domestic violence in the United States has made it a popular topic of political discourse. Housing policy advocates assert that affordable housing policies, especially rent control, are imperative for decreasing domestic violence, since affordable housing options allow victims to more easily leave their partners. However, rent control policies that include "vacancy decontrol", where the rental price can be increased whenever tenancy changes, can inadvertently lead to housing lock, a situation in which a victim stays with an abusive partner longer to avoid paying the relatively higher rent that they would face if they moved out.

In this paper, we explore the effects of San Francisco's 1994 rent control policy expansion using a continuous exposure event study design. This expansion affected only small, owner occupied buildings built before 1980, which we leverage to create zip code level measures of treatment by the policy change. We find that rent control decreased hospitalizations of women from assault. The decrease cannot be explained by changes in zip code demographics or by changes in other violent crime. These effects are consistent with the strain model of domestic violence, where financial (or other) strain in a relationship leads to violence.

The existing economic literature typically models domestic violence as a product of the bargaining power in a relationship (Aizer (2010), Munyo and Rossi (2015), Calvi and Keskar (2021), Hidrobo et al. (2016), Brassiolo (2016)) or as a result of some kind of strain or loss of control (Card and Dahl (2011) Cesur and Sabia (2016))[1]. Both of these potential channels exist in the context of housing costs. In the bargaining model, relative prices of housing in and out of the relationship will affect the level of violence within the relationship. In the financial strain model, higher housing costs will increase financial strain, leading to higher levels of violence. These models will make different predictions for the effects of specific housing policies, in particular rent control policies, on domestic violence.

A common feature of rent control and stabilization policies, including San Francisco's, is that landlords can "reset" the rent only when the tenant moves, and tenants cannot transfer the rent control provisions to their next apartment. This feature incentivizes tenants to stay in the same unit once they are under rent control, since moving would mean facing a rent increase to the market rate. These policies were structured to keep long-term residents in their neighborhoods.

However, women who are the victims of relationship violence may continue to reside with their abusive partner because of difficulty finding or affording a new apartment. In

[1]Other models of domestic violence, such as instrumental violence and male backlash exist, but do not translate to the housing context. The bargaining model shares many similarities with the commitment model of relationships used in the sociology literature (Manning et al. (2018), Kenney and McLanahan (2006), Stanley and Markman (1992))

the bargaining model of IPV, the fact that the relative cost of leaving a relationship increases with rent control would lead to a decrease in bargaining power for the victim. This decrease in bargaining power would then result in increased levels of violence. Conversely, rent control may reduce financial strain and decrease IPV triggered by financial stress. Thus, the effects of rent control on domestic violence are ambiguous.

We use data from the San Francisco Assessor's Secure Housing Roll to measure rent control exposure at the zip code level. Exposure is measured by the number of apartments with between two and four units that were built in 1979 or earlier. While we do not have data on owner occupancy as of 1994, we do have data on 1999 owner occupancy, which we use to construct an alternative measure of treatment: the number of owner occupied buildings with between two and four apartment units in each zip code.

To measure instances of domestic violence, we follow Aizer (2010), counting zip code level hospitalizations for assaults on women. We believe this to be a good proxy of domestic violence, since 82% of domestic violence victims are women (Truman and Morgan, 2014). This measure of domestic violence does not require the victim to have reported her assault to law enforcement, overcoming a common measurement issue in the crime and IPV literature. The measure also captures a very severe form of violence, since the individual will only appear in the hospital data if their injury was severe enough to lead to hospital admission.

We use a continuous treatment difference-in-differences design to determine the impact of the rent control expansion on hospitalized assaults on women. We compare neighborhoods with higher levels of newly rent controlled units to those with lower levels of newly rent controlled units before and after the policy referendum passed in 1994. We log both assaults and policy exposure variables in order to estimate an elasticity. To address the fact that there are zip codes with zero assaults in some years, we implement these specifications using the inverse hyperbolic sine.

We find that for every one percent increase in exposure to rent control in a zip code, hospitalized assaults on women decline by 0.08 percent. In levels, this translates to an almost 10 percent decrease in domestic violence for the average zip code. These results are robust to different measures of policy exposure that condition on owner occupancy, and to various alternative regression specifications. Additionally, these results are not present when we examine violence against men, suggesting that they are not due to changes in the overall levels of violence in a neighborhood. There is no strong evidence that zip code composition shifted due to the policy change, suggesting our results are not driven by changes in neighborhood demographics. We therefore believe the results reflect a change in an individual's propensity to be a victim of domestic violence.

Our effect sizes are in line with other findings in the domestic violence literature. Card and Dahl (2011) find that unexpected sports game losses can increase instances of domestic violence by 10%, Bobonis et al. (2013) find that transfers associated with the Oportunidades program decrease domestic violence by 40%, and Brassiolo (2016) finds that Spanish divorce law reform that makes it easier for individuals to divorce decreases domestic violence by 30%.

We further benchmark our findings to the overall decline in IPV in the 1990s. IPV fell from by about 60% (Rennison, 2003). Assuming San Francisco's domestic violence decrease in the 1990s followed the national trend, our result of a decrease of roughly

10% would account for roughly 16% of the decline in intimate partner violence in San Francisco.

Our paper contributes to the research on domestic violence. Many papers explore policies and phenomena that lead to changes in IPV, such as the gender wage gap (Aizer, 2010), unilateral divorce law (Stevenson and Wolfers, 2006), and transfer programs (Bobonis et al. (2013), Angelucci (2008)). We add another policy to this list, showing that rent control can decrease IPV. While most papers in this literature focus on the bargaining model of intimate partner violence, our results are more consistent with financial strain leading to violence, showing that there is not one single model that can explain intimate partner violence in every context.

We also contribute to the literature on the economics of rent control by determining the effects of rent control on the novel outcome of IPV. Previous work on the effects of rent control policies has focused on neighborhood spillovers (Autor et al., 2014a), misallocation (Glaeser and Luttmer (2003a), Olsen (1972a), and unemployment (Svarer et al., 2005). Of particular relevance are Diamond et al. (2019a), which uses the same institutional setting and policy change in 1994, and Autor et al. (2019), which studies the effects of removing rent control in Cambridge, Massachusetts on crime. Diamond et al. (2019a)track mobility at the individual level and find that rent control decreases displacement from neighborhoods, decreases mobility, and reduces the rental stock. The studied policy change in Autor et al. (2019) takes place in a similar time-frame (mid 1990s), but is a removal of rent control rather than expansion. They find a 16% decline in crime following the removal of rent control. This result contrasts with ours, where we find a decline in a specific type of crime that does not seem to be driven by a change in violent crime overall.

However, there are a couple of key differences between our studies. First, it is possible that the short run effects of removing rent control and expanding rent control are not symmetric. One possible reason for this is that it takes time for the market rent to diverge from the rent controlled rent, but once these differences have emerged, removing rent control collapses them immediately. A second key difference between this setting and ours is that Cambridge did not allow for vacancy decontrol, where landlords are allowed to re-set the rent to market between tenants. This means that the value of a rent controlled unit does not increase in the length of time that you remain in the unit. Finally, we are limited in our hospital data to looking only at violent crimes that result in severe bodily injury; it is possible that there were effects on property crime or more minor violent crimes that we are unable to detect.

Finally, this project is related to broader research on the social implications of a lack of affordable housing. While rent control is a particularly strong example of housing-related incentives to stay in an abusive relationship, we might expect to see similar effects in tight housing markets, where it would be difficult to find an apartment, afford an apartment, or break a lease mid-term.

The rest of this paper is organized as follows. In Section 2.2, we discuss the financial strain and bargaining models of domestic violence and how rent control fits in to both of these models. In Section 2.3, we discuss rent control policy details for San Francisco and the policy change that we exploit as a natural experiment. In Section 2.4, we provide details of how we construct measures at the zip code level of exposure to the policy change and IPV. In Section 2.5, we explain our empirical strategy, the required assumptions, and the robustness checks we perform. In Section 2.6 , we present results, then in Section 2.7 we explore alternative mechanisms. Finally, we conclude in Section 2.8.

2.2. Conceptual Framework

We consider two primary models of domestic violence: the financial strain model and bargaining model. These models make different predictions of the consequences of a rent control policy. Both models incorporate the fact that rent control effectively increases the household budget of couples: the financial strain model focuses primarily explores that expanded budget set, while the bargaining model focuses on how that budget set is now comparatively larger relative to the budget set outside the relationship, changing the bargaining dynamics in the relationship.

2.2.1. Financial Strain Model

In a financial strain model, loosely based on Card and Dahl's (2011) loss-of-control model, there is some probability of violence breaking out in the relationship for any given time period. This probability of violence depends on the level of underlying stress in the relationship, where one large component is the level of financial strain. As financial strain increases, so does the likelihood of experiencing violence in the relationship.

Rent control changes the level of financial strain by shifting out the budget set for households that are the beneficiaries of rent control. Rent controlled tenants pay relatively less rent than their non-rent controlled peers over time: as market rent prices increase, non-rent controlled tenants' rent increases, but rents stay the same for rent controlled units. This relative decline in rent prices leads to a relative decline in financial strain, resulting in less violence. In the long run, these effects may attenuate if general equilibrium forces increase rental prices for non-rent controlled units in the area. In the financial strain model, we would expect to see a decline in IPV due to rent control.

The financial strain model has been shown to drive violent behavior in a number of contexts. Cesur and Sabia (2016) use the strain theory to explain why veterans who were engaged in active duty were more likely to commit intimate partner violence and child abuse. Card and Dahl (2011)'s results also align with the strain model: when an individual's local sport team unexpectedly loses, the strain of that loss leads to an increase in intimate partner violence. This phenomenon is not unique to housing and intimate partner violence: Holz et al. (2020) show that stress due to a peer's injury leads police officers to act more violently in their use of force.

2.2.2. Bargaining Model

An alternative model of IPV is the bargaining model. Here, the abuser values both consumption and violence, and his partner values consumption and safety. The couple then bargains over the level of violence in the relationship, where the outside option of leaving the relationship involves the victim moving out. When the outside option becomes less attractive to the victim, it becomes more difficult to leave the abusive partner, which then decreases the victim's bargaining power and leads to increased violence. Under rent control, the victim's outside option declines relative to the abuser's outside option of staying in a rent controlled unit, since for the victim, leaving would mean paying the relatively higher market-rate rent.

The bargaining model has been used in numerous economics papers to explain intimate partner violence patterns. These include Aizer (2010), Munyo and Rossi (2015), Calvi and Keskar (2021), Hidrobo et al. (2016), and Brassiolo (2016). The bargaining model also maps conceptually to one of the main models of domestic violence discussed in the sociology literature, the commitment theory model. This model is discussed in relation to housing and cohabitation. In this model, the level of instability and violence in a relationship is directly related to the commitment level in a relationship. When couples move in together, constraints (such as housing constraints) increase without a corresponding increase in commitment, as would be common in a marriage (Manning et al. (2018), Rhoades et al. (2010), Kenney and McLanahan (2006), Stanley and Markman (1992)). The increase in constraints without commitment can lead to violence. Additionally, whether a woman stays with her abuser has been linked to the level of resources she has access to and the degree of power within the relationship (Gelles, 1976). In a bargaining model, any factor that increases difficulty in leaving a relationship is considered to decrease bargaining power of the partner who wants to leave. In this framework, marriage, or any analogous increase in commitment, could be thought of as changing the bargaining problem as whole.

Under the strain model of domestic violence, we would expect to see a decline in intimate partner violence, whereas under the bargaining framework, we would expect to see an increase. Since other models of domestic violence do not make sense within our context, we can infer which model more accurately characterizes behavior based on the direction of our results. A decline in intimate partner violence in response to the rent control expansion would be consistent with the financial strain model of domestic violence, but not the bargaining model.

2.3. Rent Control in San Francisco

In 1980, San Francisco passed its Rent Ordinance. This enacted a rent control policy for existing buildings with five or more units and smaller buildings with units that were owner-occupied. Landlords in these buildings were not allowed to raise rents by more than a statutorily established amount and must renew leases. The policy did not extend to new construction; only buildings built in 1979 or earlier were controlled. Owner-occupied buildings with four or fewer units were exempted from the policy to protect "mom-and-pop" landlords.

In November 1994, the city passed by referendum a new rent control law, Proposition I, which lifted the exemption on small buildings, effectively controlling all buildings built before 1980. Units built after 1980 are never subject to rent control.

Most landlords of newly controlled buildings were not allowed to raise rents by more than 2% (an upper bound which had been decreased by half in 1992 under Proposition H) until the current tenants moved out. At that point, the landlords could reset the rent of the apartment to the market rate, with some exceptions for landlords who had not historically raised rent prices.[2] This policy change affected zip codes across San Francisco

[2]For example, if a landlord had not increased the price between 1991 and 1994, they were entitled to raise the rent by 7.2%, as long as they filed a petition with the city to do so and gave the tenant proper notice of the increase. The longer the period of no increase, the higher the landlord was entitled to raise prices. The largest allowed increase was 15.2% for landlords who had not raised rent between 1989 and 1994. The ordinance also required that rent increases made after May 1st of 1994 be refunded to tenants, possibly leading to large lump-sum payments to tenants in newly controlled units. We do not have any data on whether such payments occurred in practice.

variably, depending on how many small (four units or fewer), old (built in 1979 or before), owner-occupied buildings exist in the zip code.

San Francisco's rent control policy allows landlords to raise the rent in an unlimited fashion if a tenant leaves of their own volition. This feature contrasts with rent stabilization policies found elsewhere (Los Angeles, Cambridge, New York for tenants who moved into units after 1971) that restrict the rent of a unit, regardless of whether the tenant moves or stays. These policies may have different impacts on domestic violence than San Francisco's rent control policy. Recent rent control policies, such as that passed in Oregon, share this feature with San Francisco, where landlords are able to reset rent to market in between tenants, but are limited in how much they can raise rent for a given tenant.

2.4. Data

We use data from two main sources. Data on the number of newly rent controlled units comes from the San Francisco Assessor's Office. Data on the number of hospitalizations resulting from assaults come from California's Office of Statewide Health Planning and Development (OSHPD) from 1990-2000. We supplement these data with information from the 1990 and 2000 Census on zip code level characteristics.

2.4.1. Measuring Rent Control

We use data from the San Francisco Assessor's Secure Housing Roll from 1999 to determine the number of units treated by the policy change at the zip code level. These data include the owner's mailing address, the address of the property, and the year the building was

Figure 2.1. Zip Code Level Exposure to Change in Rent Control Policy

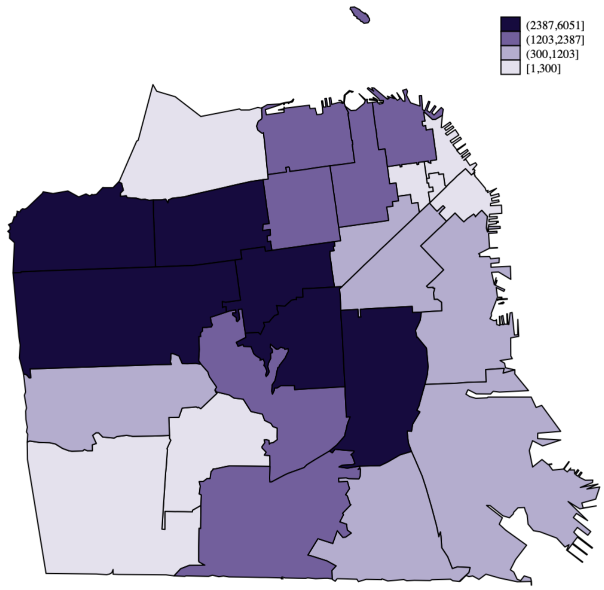

Notes: This map depicts exposure at the zip code level for the San Francisco policy change. Exposure is measured as the number of units in a zip code in buildings with 2-4 units that were built prior to 1979. Data source: 1999 San Fransisco Assessor's Secure Housing Roll and authors' calculations.

built, and the number of units in the building. We restrict to building with residential use codes. Appendix B.1 discusses how we handle missing unit numbers, building ages, or zip codes. To validate our data cleaning choices, we compare our final measures of the housing stock against measures from the Census in Appendix Table B.1. There are only minor discrepancies that could be explained by the addition or demolition of buildings between 1999 and 2000.

We identify units as treated if they are located in buildings built prior to 1980 with two to four units and aggregate to the zip code level.[3] Figure 2.1 shows this measure of treatment plotted on a map of San Francisco.

[3]The most granual geographic information on hospitalizations is at the zip code level.

We additionally construct several alternative measures of treatment, designed to account for potential mismeasurement due to the fact that the earliest version of the Assessor data we could obtain is from 1999, several years after the policy. First, we develop a version that accounts for owner occupancy. We define owner occupancy as having matching addresses (street number and street name), where the city and state of the mailing address are San Francisco, CA. Second, we attempt to account for the fact that buildings treated by rent control could have been converted into condos or demolished and replaced in response to the policy. We do so by varying how we classify condos and new construction to create various alternative measures of treatment.

Figure B.1 shows these alternative measures of treatment. Table B.2 reports the correlation matrix betwee these various measures. These measures of treatment are all highly correlated, largely because there was not substantial new construction in San Francisco between 1995 and 1999.

2.4.2. Measuring Intimate Partner Violence

To measure intimate partner violence, we use administrative data on hospital admissions from California's Office of Statewide Health Planning and Development (OSHPD) to determine the number of hospitalizations resulting from assault. These data cover the universe of hospital admissions in California's hospitals and contain information on patient zip codes[4], patient demographics, and diagnostic codes. We use the External Cause of

[4]We drop some observations due to missing zip codes. 0.6% of the observations are dropped because the zip code was labeled as "unknown", 0.1% of the observations are dropped because the patient lives outside of the US, and 2.1% of observations are dropped because the patient is homeless.

Injury Codes (E-codes), which state the underlying cause of the injury that resulted in the admission.

These diagnostic codes have been previously used in the economics literature on domestic violence (Aizer, 2010). Aizer (2010) reports that about 80% of assaults on women are related to IPV. We focus primarily on women since 82% of intimate partner violence is committed against women (Truman and Morgan, 2014). We additionally exclude children from our sample.

We identify a patient as having been assaulted if they are assigned E-codes describing an assault as the reason for the hospital admission. We list the specific E-codes that we use and their definitions in Appendix B.2. This measure of intimate partner violence does not require reporting to the police, so it avoids some of the drawbacks typically associated with measures of domestic violence constructed from criminal records. A health practitioner only needs to suspect an injury is due to assault in order for it to show up as possible IPV in our dataset.

One caveat to this measure is that it measures some of the most severe instances of domestic violence. Roughly half of IPV results in injuries, with 11% of incidents resulting in serious injuries (Truman and Morgan 2014). About 34% of women who are injured in IPV need medical care, but only half of women who need medical care seek it in a professional medical setting, as opposed to seeking care from a neighbor or family member (Truman and Morgan, 2014). Based on a back-of-the-envelope calculation using these statistics, our measure of domestic violence will capture roughly the most severe 8% of the instances of intimate partner violence.

Another important characteristic of this measure of IPV is that it may capture a change in the intensity rather than the prevalence of violence. It is unclear if violence significant enough to result in hospitalization scales linearly with other forms of violence, so the interpretation of our results must be cautious. A reduction of 10% in hospitalizations may result from a 10% reduction in overall violence; however, it also could be caused by a decrease in the severity of violence, with total instances of violence remaining constant. While we may not be able to distinguish between these two, we note that reducing both severity and quantity of violence is beneficial.

We also identify men who have been hospitalized as a result of an assault, which we can use to control for underlying trends in violence. While some of these men's assaults may have resulted from IPV, the bulk of the assaults likely were not. Only 10% of assaults on men can be attributed to domestic violence (Truman and Morgan, 2014).

Our measure of intimate partner violence conforms to known patterns in domestic violence. Based on data from the National Crime Victimization Survey, domestic violence most commonly affects women between the ages of 18 and 24 and Black women of all ages (Truman and Morgan, 2014). We show in Appendix Figures D1 and D2, which plot the distribution of assaulted patient characteristics and the characteristics of the full patient population, that we match these two fact patterns.

Appendix Figure C5 shows the time patterns in our measure of IPV, split by quartiles of treatment exposure. IPV is flat over time in the lowest treatment quartile. It falls over time in the higher three treatment quartiles.

Table 2.1. Summary Statistics

Panel A: Characteristics	(1) Full Sample	(2) Low Treatment	(3) High Treatment	(4) Difference
Median Income 1990	45,698.68	44,922.08	46,415.54	1,493.46
Median Income 2000	58,886.84	56,642.83	60,958.23	4,315.40
Median Rent 1990	869.28	869.00	869.54	0.54
Median Rent 2000	1,006.36	1,028.83	985.62	-43.22
Pre-Policy Black Patients	495.28	511.77	480.06	-31.71
Post-Policy Black Patients	435.45	472.61	401.14	-71.47
Pre-Policy White Patients	1,642.54	968.77	2,264.48	1,295.71***
Post-Policy White Patients	1,640.53	972.31	2,257.36	1,285.05***
Pre-Policy Asian Patients	531.16	352.22	696.34	344.122**
Post-Policy Asian Patients	607.56	419.93	780.76	360.826*
Pre-Policy Hispanic Patients	322.14	130.27	499.246	368.979
Post-Policy Hispanic Patients	298.37	126.24	457.256	331.02
Pre-Policy Median Age	54.00	52.17	55.692	3.526
Post-Policy Median Age	58.87	56.90	60.686	3.79
Pre-Policy Patients	3,043.90	1,997.72	4,009.62	2,011.89***
Post-Policy Patients	2,870.79	1,957.99	3,713.37	1,755.39**
Total Housing Units	12,622.28	6,927.42	17,879.08	10,951.66***
Panel B: Treatments and Outcomes				
Pre-Policy Assaults on Women	7.74	6.98	8.45	1.46
Post-Policy Assaults on Women	4.21	3.71	4.68	0.97
Pre-Policy Assaults on Men	37.81	30.48	44.57	14.09
Post-Policy Assaults on Men	18.69	15.69	21.45	5.75
# Prev. Rent Controlled	4,882.08	2,412.17	7,162.00	4,749.83**
# Treated	1,688.48	358.42	2,916.23	2,557.81***
Observations	25	12	13	25

Notes: This table reports summary statistics split by the level of treatment of the zip code. The full sample includes all zip codes. Low treatment zip codes have fewer than the median number of units that became newly rent controlled in 1994. High treatment zip codes have more than the median number of units that became newly rent controlled in 1994. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

2.4.3. Demographic Characteristics

We measure the demographic characteristics of zip codes in the 1990 and 2000 Census and from the demographic characteristics of hospital admissions. Panel A Table 2.1 reports averages of these characteristics both before and after the policy change. Column (1) reports averages for the entire sample of 25 zip codes. Columns (2) and (3) divide the sample into zip codes whose level of treatment was below the median level of treatment and zip codes whose level of was above the median. Column (4) reports the difference between the two columns and whether the differences are statistically significant.

We find that the low and high treatment zip codes are similar on non-population related characteristics both before and after the policy. These zip codes have comparable incomes, rents, patient ages, and minority patients. High treatment zip codes have more patients, more White patients, and more total housing units.

Panel B reports treatment and outcome averages for these three groups of zip codes. We find that high treatment zip codes have more assaults on women pre-policy, although this difference is not statistically significant. There is a smaller difference, although also not statistically significant, post policy. This pattern is similar for assaults on men. By construction, high treatment zip codes have more newly rent controlled units. They also have more previously rent controlled units.

In Appendix Table B.4, we assess whether Census characteristics can be used to predict the number of treated units. We do not find a statistically significant relationship between any Census characteristics and the number of treated units, and the R^2 of this regression is quite low.

2.5. Empirical Strategy

We use a difference-in-differences strategy to determine the effect of rent control on IPV. We compare zip codes with high policy exposure to zip codes with low exposure before and after the 1994 policy change. We measure exposure using the number of small apartments built before 1980 in each zip code.

We estimate the following specification:

$$\text{Log}(\#\text{ assaults on women}_{it}) = \alpha_i + \tau_t + \beta \cdot \text{Log Exposure}_i \times \text{Post}_t + \epsilon_{it} \tag{2.1}$$

where α_i are zip code level fixed effects, τ_t are year level fixed effects, Exposure_i is the number of units newly treated by the rent control expansion in zip code i, and $Post_t$ is an indicator for years after 1994.

We estimate corresponding specifications in an event study framework, using the following specification:

$$\text{Log}(\#\text{ assaults on women}_{it}) = \alpha_i + \tau_t + \sum_{\tau=1990,\tau\neq 1994}^{2000} \beta_\tau \cdot \text{Log Exposure}_i \cdot \mathbf{1}\{\text{Year} = \tau\} + \epsilon_{it} \tag{2.2}$$

where we estimate different coefficients β_τ for the interaction of our exposure variable with each year leading up to and after the policy change. We omit 1994, which is the year the referendum passed, as our reference year. This event study specification allows us to visually assess for the presence of pre-trends and assess whether there is a time pattern in the response to the policy.

We additionally estimate specifications that allow zip codes in different terciles of treatments to be on different linear time trends. We do a similar exercise for terciles of assaults against men in the pre-period, which allows zip codes with different baseline violence rates to be trending differently. We also include controls for the number of previously rent controlled units and for the demographics of women admitted to San Francisco hospitals from give zip codes. These address concerns that rent control causes equilibrium changes in the housing market and in neighborhood demographics, respectively.

Interpretation of a causal effect requires four assumptions. First, we assume that there were no anticipatory effects of the policy. The passage of this policy was unexpected; the policy passed in a close election, receiving 51% of votes (Harrison, 1994). It is thus unlikely that there would be changes in household behavior in anticipation of the policy.

Second, we assume that absent treatment, zip codes with high exposure to the rent control policy would have trended similarly to those with low levels of exposure. We can examine pre-trends estimated in an event study specification for evidence that this assumption is violated.

Third, we must assume that there are no spillovers across units that were differentially treated. This assumption is difficult to test and may be a strong assumption in our settings; if the rent control policy change alters displacement across zip codes, less treated zip codes may be affected by the policy in subtle ways. There are two drivers of spillovers that we are concerned about. The first is that the expansion of rent control may cause changes in the broader market. To address this, we can estimate alternative specifications that include the number of previously rent controlled units as zip codes with more pre-existing rent controlled units may be differentially affected by spillovers. We also use Census data to look at the effects on median rent prices. The second is that new residents to San Francisco may have chosen to locate in highly treated neighborhoods in the pre-period and must locate in less treated neighborhoods in the post-period. We address this by looking at the effects of the policy on neighborhood demographics.

Finally, we assume that there is homogeneity in treatment effects, meaning that potential treatment effects must be unrelated to policy exposure (Callaway and Sant'Anna, 2021). This assumption is required due to the continuous nature of our specification. If individuals that are at higher risk for intimate partner violence (and therefore could have larger treatment effects if treated) all live in neighborhoods that happened to be heavily treated by the policy change, then this assumption would not be met. Because we cannot accurately predict the number of treated units using Census demographics and do not find large differences across zip codes by treatment levels except through variables related to the size of the zip code, we think this assumption is reasonable.

While per-capita estimates would be desirable, the data required to estimate these does not exist. In the 1990s, there are not reliable zip code-level estimates of population at the annual level. If instead we were to estimate policy exposure as a proportion of buildings in a zip code, there would be no analogous denominator for the outcome of assaults, which should also scale with population. For these reasons, we instead focus on estimating elasticities in order to explore proportional effects.

To account for the fact that some zip codes experience no assaults in a given year, we implement these specifications using the inverse hyperbolic sine. We test robustness to this choice by using the more standard log specification and dropping observations with zero assaults and by using $log(1 + assaults)$ instead. Verifying that our estimates are not driven by choice of model specification is important in this setting because the inverse hyperbolic sine may be a less accurate approximation to logged values when it is taken over small values (Bellemare and Wichman, 2020), and the number of assaults on women is low. We additionally estimate a Poisson specification since a count model of assaults may be appropriate in our setting. Finally, we estimate alternative specifications in levels.

To test for alternative explanations, we estimate specifications similar to Equation 2.1 with male assault hospitalizations as well as the number of hospitalizations for various demographic groups on the left hand side.

2.6. Effect of Rent Control on Female Assault Hospitalizations

We begin by estimating Equation 2.2 for female assault hospitalizations. As shown in Figure 2.2, there is a decrease in hospitalized assaults on women following the implementation of the policy. Here, we plot the coefficients of β_τ, the coefficient on the number of treated units interacted with dummy variables for each year. The purple dashed line shows the difference-in-difference estimate of β in Equation 2.1.

This plot presents visual evidence in support of our assumption of parallel trends. There is no evidence that zip codes with more treated units were trending differently in the period before the policy chage.

In Table 2.2, we report the difference-in-difference estimate shown in the purple dashed line, along with estimates from alternative specifications. Column 2 allows zip codes who are in different terciles of treatment to be on different time paths by including tercile-specific linear time trends. Column 3 includes male assault tercile-specific linear time trends. Column 4 controls for the inverse hyperbolic sine of the number of previously rent controlled units. Column 5 includes controls for admitted patient characteristics.

Figure 2.2. Effects of Rent Control Exposure on Female Assault Hospitalizations

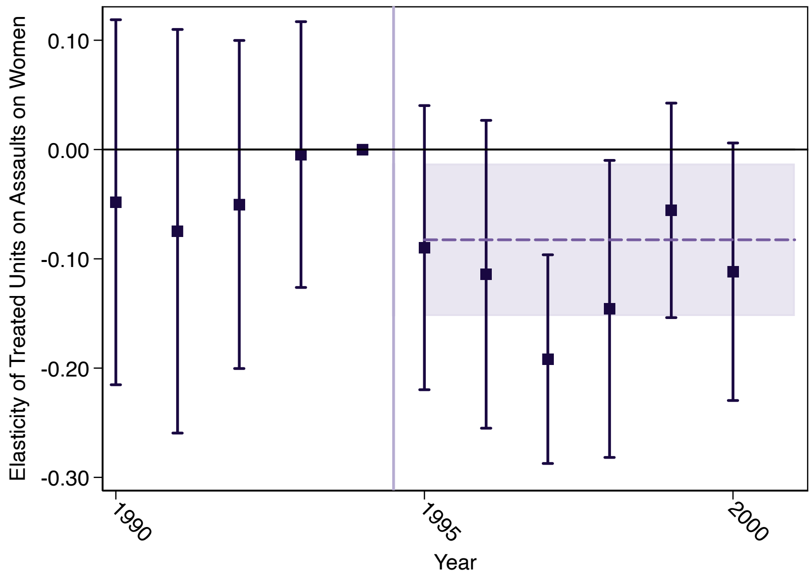

Notes: This figure shows our estimates of the event study specification in equation 2.2. Each point corresponds to the coefficient β_τ. This regression uses the inverse hyperbolic sine of the number of units as the treatment variable and the inverse hyperbolic sine number of hospitalizations per zip code as the outcome. Error bar show 95% confidence intervals with standard errors clustered at the zip code level. The dashed line shows the estimated difference-in-differences coefficient. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

We find that a 1% increase in the number of newly rent controlled units results in a 0.08% decrease in the number of assaults resulting from domestic violence. The average number of newly rent controlled units is 1,688 and the average number of assault pre-policy is 7.7, so this is equivalent to 16 additionally rent controlled units leading to a reduction of .006 assaults. These results translate to just under a 10% decrease in IPV for the average zip code. This reflects a meaningful change in the number of assaults, given that we are capturing a very severe form of intimate partner violence.

Table 2.2. Effects of Rent Control Exposure on Female Assault Hospitalizations

	(1) Assaults	(2) Assaults	(3) Assaults	(4) Assaults	(5) Assaults
Post=1 × IHS NumTreated	-0.0826**	-0.105***	-0.0837**	-0.0684	-0.111*
	(0.0338)	(0.0316)	(0.0391)	(0.0795)	(0.0627)
Treatment Tercile Specific Trends		X			
Male Assault Tercile Specific Trends			X		
Previously Rent Controlled				X	
Hospital Covariates					X
R-Squared	0.812	0.816	0.813	0.812	0.818
Dep Var Mean	5.82	5.82	5.82	5.82	5.82
Observations	275	275	275	275	275

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the results of regressions of the IHS of assaults on women on a post-1994 indicator interacted with the IHS of the number of treated units. Column 1 reports the results of equation 2.1. Column 2 includes treatment tercile specific linear time trends. Column 3 includes male assault tercile linear time trends. Column 4 includes a control for the IHS of the number of units that were previously rent controlled in a zip code. Column 5 includes hospital patient characteristics controls. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

Our estimated effects are robust across the alternative controls we include. Including the number of units that were previously rent controlled attenuates the result somewhat, but largely serves to make our estimates noisier. Appendix Figure C1 shows the estimates from the equivalent event study specifications.

In Table 2.3, we report estimates from alternative ways of specifying our regression model. Recall that we have two issues we want to address in our choice of specification: it is not possible to calculate per capita assaults and there are several zip codes with zero assaults. For comparison, Column 1 again reports our preferred estimates using an inverse hyperbolic sine. Column 2 is a log-log specification. This specification drops observations for zip codes with zero assault hospitalizations. Column 3 addresses this issue in an alternative way by adding 1 to the number of assaults for each zip code. Column 4 is a Poisson model. Column 5 reports our estimates in levels.

Our results are consistent across different ways of implementing a log specification. The effect sizes are attenuated in the Poisson specification. Our estimates are noisier in the level specification, but the effect sizes are consistent with what we find in our preferred specification. Appendix Figure C3 shows the estimates from the event study specifications for these alternative specifications.

Table 2.3. Alternative Specifications

	(1) IHS Assaults	(2) Log Assaults	(3) Log(1+Assaults)	(4) Log(Assaults)	(5) Assaults
Post=1 × IHS NumTreated	-0.0826**				
	(0.0338)				
Post=1 × Log NumTreated		-0.100*		-0.0471	
		(0.0493)		(0.0567)	
Post=1 × Log(1+NumTreated)			-0.0729**		
			(0.0285)		
Post=1 × NumTreated					-0.390
					(0.464)
Specification	IHS	Log	Log(1+X)	Poisson-Log	Levels
R-Squared	0.812	0.799	0.821		0.780
Dep Var Mean	5.82	5.82	5.82	5.82	5.82
Observations	275	233	275	275	275

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports results of alternative regression specifications. Column 1 is our preferred specification. Column 2 is a log-log specification. Column 3 is log(1 + x) on both sides. Column 4 is a Poisson specification. Column 5 reports our effects in levels. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

We additionally estimate specifications that address the noise that is present in our measurement of treatment. We account for owner occupancy and address that the policy may have led to condo conversions or new construction in several ways. We show estimates from specifications that include these alternative measures of treatment in Appendix Table B.3 and Figure B.3. Perhaps not surprisingly, given the high correlation between our preferred measure of treatment and these alternative measures, our estimates of the effects of rent control are largely similar.

The resulting decrease in domestic violence falls well within the range of treatment effects found by other papers in the literature. Card and Dahl (2011) find that unexpected sports game losses can increase instances of domestic violence by 10%, Bobonis et al. (2013) find that transfers associated with the Oportunidades program decrease domestic violence by 40%, and Brassiolo (2016) finds that Spanish divorce law reform that makes it easier for individuals to divorce decreases domestic violence by 30%.

IPV fell nationally by 60% over the course of the 1990s (Rennison, 2003). If San Francisco followed this trend, then our results suggest that rent control accounted for about 16% of the decline in intimate partner violence in the 1990s.

2.7. Mechanisms

Thus far, we have shown that the expansion of rent control in San Francisco led to decreases in the number of hospitalizations due to assault of women in areas that were heavily affected by the policy. However, it is possible that these decreases were driven by factors other than IPV. For instance, if the policy change resulted in a decrease in either crime overall or the propensity to seek medical care, then it is possible that our results are being driven by those factors, rather than by a decrease in IPV. Additionally, it is possible that our effects could be driven by demographic changes in who lives in areas that are heavily rent controlled. In this section, we present evidence consistent with our effects being primarily driven by IPV, rather than other factors.

2.7.1. Violence

One possible cause of our results that does not involve changes in intimate relationships is a change in the level of violent crime. To test whether the the decrease in assaults on women we estimate is due to a decrease in violent crime overall, we estimate the effects on assaults on men, who are less likely to be assaulted due to IPV and more likely to be assaulted due to general violent crime. Truman and Morgan (2014) document using the National Crime Victimization Survey that 55% of serious violent crime against males is committed by strangers.

In Figure 2.3, we find a null effect of exposure on assaults resulting in hospitalization for men, suggesting the decreases in assaults on women are not due to changes in violent crime overall. In later years, there are negative estimates of the effects of the policy; these effects are not statistically significant and are not consistent with the years in which we see the largest effects for women, which are the years immediately following the policy chage.

2.7.2. Hospitalizations

Given that rent control alters the budget constraint of households who are the beneficiaries, one concern is that individuals who are the beneficiaries may be more likely to seek

Figure 2.3. Effects of Rent Control Exposure on Male Assault Hospitalizations

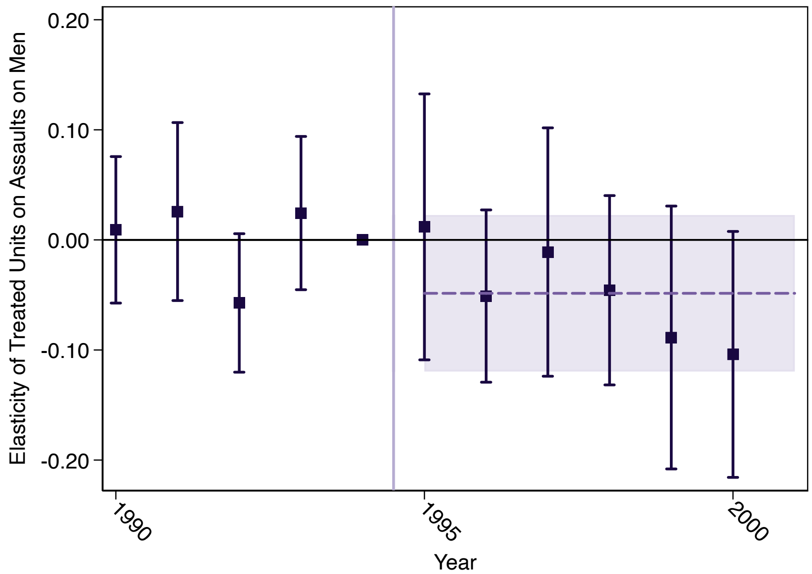

Notes: This figure shows our estimates of the event study specification in equation 2.2 with male assault hospitalizations as the outcome. Each point corresponds to the coefficient β_τ. This regression uses the inverse hyperbolic sine of the number of units as the treatment variable and the inverse hyperbolic sine number of hospitalizations per zip code as the outcome. Error bar show 95% confidence intervals with standard errors clustered at the zip code level. The dashed line shows the estimated difference-in-differences coefficient. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

medical care. If this were the case, we would expect to see an increase rather than decrease in domestic violence. We additionally can test whether this is the case by examining hospitalizations resulting from accidents, by using e-codes in a similar manner as we do with assaults. Appendix B.2 lists the specific codes that we use to identify accidents. Figure 2.4 shows the effects of rent control on accidents for both men and women in subfigures (a) and (b). We find no effects for men and positive effects for women. These effects would bias our estimates upwards if they reflect an increased tendency to seek medical care following the policy change.

We also look at the effects on drug and alcohol abuse, which could contribute to a loss of control and more violence. We show our event study estimates for these hospitalizations in subfigures (c) and (d). We find no evidence that rent control changed the number of substance abuse hospitalizations.

2.7.3. Demographic Changes

Because we do not have data at the individual address level, it is difficult to distinguish between changes in intimate partner violence that come from changes in individual behavior versus changes in who lives in a given zip code who may have different propensities for violence. Both explanations may be relevant for policy makers, but have different interpretations; the first suggests a true drop in violence, the second suggests a displacement of that violence.

To explore whether there is evidence of large demographic changes, we would ideally estimate our difference-in-difference specification with various measures of zip code level demographics, such as age, income, and race, that we know are correlated with domestic violence on the left hand side. Unfortunately, zip code level characteristics are not available at an annual level during our sample period. Instead, we use Census level characteristics from 1990 and 2000 and patient characteristics at the zip code level at an annual level and estimate whether we observe changes in these characteristics post policy.

Table 2.4 reports our estimates looking at the effects on Census level zip code characteristics. We examine the effects on population, median income, median rent, the percent

Figure 2.4. Effects of Rent Control Exposure on Other Hospitalizations

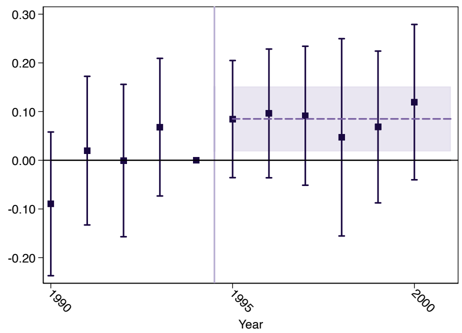

(a) Female Accident Hospitalizations

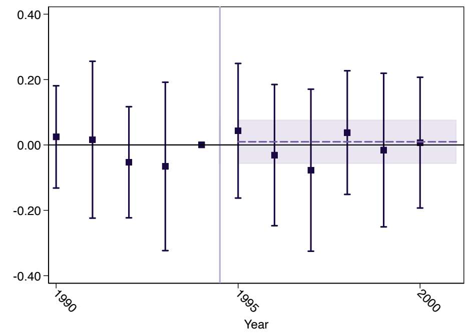

(b) Male Accident Hospitalizations

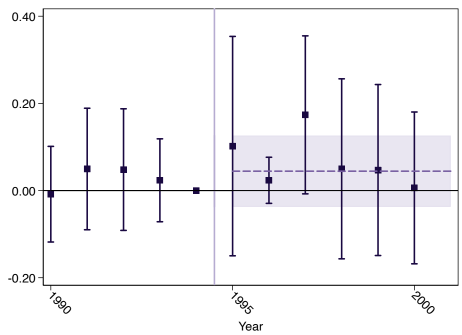

(c) Female Drug and Alcohol Hospitalizations

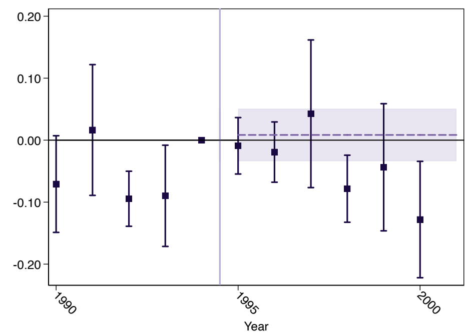

(d) Male Drug and Alcohol Hospitalizations

Notes: This figure shows our estimates of the event study specification in equation 2.2 with the alternative types of hospitalizations as the outcome. Each point corresponds to the coefficient β_τ. This regression uses the inverse hyperbolic sine of the number of units as the treatment variable and the inverse hyperbolic sine number of hospitalizations per zip code as the outcome. Error bar show 95% confidence intervals with standard errors clustered at the zip code level. The dashed line shows the estimated difference-in-differences coefficient. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

living under the federal poverty line, and the population that is female and 18-24.[5] We find small positive effects on the size of the population. Given our effects are not expressed per capita, this change would bias us towards finding increases in the number of assaults,

[5]Race and ethnic categories changed between the 1990 and 2000 Census so we do not include them here.

Table 2.4. Effects of Rent Control on Census Characteristics

	(1) IHS Population	(2) IHS Median Income	(3) IHS Median Rent	(4) IHS % Poverty	(5) IHS Female 18-24
Post=1 × IHS NumTreated	0.0666** (0.0284)	-0.00460 (0.0198)	-0.0307 (0.0224)	-0.0139*** (0.00252)	0.0393 (0.0395)
R-Squared	0.993	0.977	0.951	0.979	0.986
Dep Var Mean	29894	52293	938	.14	1398
Observations	50	50	50	50	50

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports estimates of a difference-in-difference regression on Census characteristics at the zip code level. Data Sources: 1999 San Francisco Assessor's Secure Housing Roll, US Census Bureau

rather than the declines we find. We find no effects on the median income, median rents, or the number of women 18-24. We look at the effects on the number of women 18-24 because this group is the group that is most likely to be the victims of IPV. Finding no effects here suggests that our effects are being driven by changes in behavior rather than changes in the composition of the zip code.

We do find small declines in the share of the population living under the poverty line, even though median income does not change. Given poverty is one predictor of IPV, this change could contribute to our results. It is relatively small compared to the result of our results, so we think it is unlikely to be the main driver of our results, but could contribute to the negative effects we find.

Because these characteristics are only available in 1990 and 2000, it is impossible to determine if zip codes were trending similarly in these characteristics. Rather than the

Table 2.5. Effects of Rent Control on Patient Characteristics

	(1) IHS Black Patients	(2) IHS White Patients	(3) IHS Hispanic Patients	(4) IHS Asian Patients	(5) IHS Median Age
Post=1 × IHS NumTreated	0.0511 (0.0817)	0.0310 (0.0328)	0.0605* (0.0316)	0.0356 (0.0308)	-0.0135 (0.0128)
R-Squared	0.975	0.986	0.977	0.989	0.869
Dep Var Mean	462.64	1641.44	309.18	572.84	56.66
Observations	275	275	275	275	275

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports estimates of a difference-in-difference regression on patient characteristics at the zip code level. Data Sources: OSHPD Inpatient Database, US Census Bureau

effect of the policy, these estimates may reflect the highly treated areas were trending differently before the policy in these characteristics.

For this reason, we now examine whether there are effects of the policy on patient characteristics in the inpatient data. Table 2.5 reports these estimates. We find no effects of the policy on the number of Black, White, or Asian patients or on the median age of patients. We find increases in Hispanic patients. However, contemporaneously with our policy change, the classification of race and ethnicity in our hospital data changed, which may have affected this result.

To aggregate these measures, we test in aggregate whether any of demographic changes should have lead to a change in predicted intimate partner violence, based on characteristics alone. To predict intimate partner violence, we first run a regression of IPV on the number of patients of each race who were admitted to the hospital, and median age of hospitalized patients. The results of this predictive equation are shown in Appendix Table D1.

Figure 2.6 shows our event study estimates for this predicted measure of intimate partner violence. We do not find changes in predicted assaults on women based on demographic changes at the zip code level. Because of this, we believe it is unlikely that changes in demographic characteristics of zip codes are the primary drivers of our results.

Figure 2.6. Effect of Rent Control Exposure on Predicted Female Assault Hospitalizations

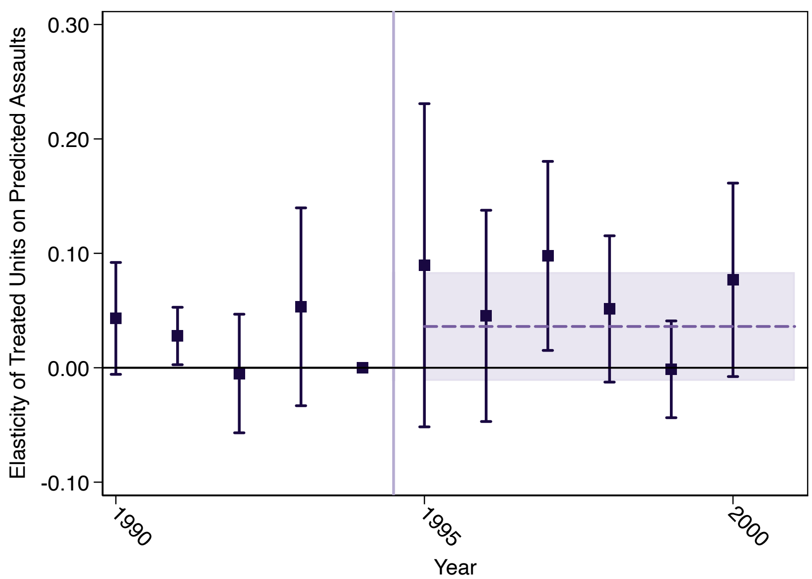

Notes: This figure shows our estimates of the event study specification in equation 2.2 with our predicted measure of female assault hospitalizations as the outcome. Each point corresponds to the coefficient β_τ. This regression uses the inverse hyperbolic sine of the number of units as the treatment variable and the inverse hyperbolic sin number of hospitalizations per zip code as the outcome. Error bar show 95% confidence intervals with standard errors clustered at the zip code level. The dashed line shows the estimated difference-in-differences coefficient. Data sources: 1999 San Francisco Assessor's Secure Housing Roll, OSHPD Inpatient Database

2.7.4. Housing Market Changes

It is unclear what the aggregate effects of rent control on rent were. While rent control mechanically lowers rent for those who receive the benefits of rent control, it may have the effect of raising rents for the rest of the market as it limits the supply of market rent housing. Unfortunately, records of zip code level rents from the 1990s are difficult to obtain. To assess the effects of rent control on rent, we use data on median gross rents from the US Census. These rents are self-reported by individuals who respond to the long form version of the Census. We only have three years of data available to use: 1980, 1990, and 2000, which makes it difficult to evaluate how rents were evolving prior to the passage of the 1994 referendum or the time pattern of rent prices after its passage.

We estimate a similar event study specification to our main specification with these three years of rent data and do not find strong evidence of rent declines following the policy. The event study plot is available in Appendix Figure H1.

This analysis likely masks heterogeneity in the effects on rent prices. Rent control will have the biggest effects on units where market rents are rising, so while rents may be staying low for some units in a zip code, they may be simultaneously rising for other units in an area, creating a wedge between rent controlled rents and market rents.

There are likely other changes in the housing market as well. Diamond et al. (2019a) find that rent control limits renters' mobility but decreases the rental supply as landlords redevelop buildings or convert to condos. In a companion working paper, we find that the expansion of rent control increases the number of wrongful eviction claims, likely related to improperly done owner-move in or Ellis Act evictions. These changes may lessen the benefits of rent control in the medium to long term, as renters whose units are removed from the rental stock no longer have the benefits of rent control, consistent with the fact that we see a gradual fade-out of some of our effects over time.

2.8. Conclusion

In this paper, we study the relationship between housing prices and intimate partner violence. There are two major models of intimate partner violence that are applicable in the setting of housing prices: the financial strain model and bargaining model. These models make different theoretical predictions in the case of rent control, where there exists a stark difference between the rent prices made in a relationship and outside of the relationship. Determining which channel dominates is important when thinking about housing affordability policy, given that helping victims of domestic violence is one reason (among many) that advocates push for new policies to improve affordability. Understanding whether violence related to housing is the result of stress in the relationship or about a victim's ability to leave and their outside option will guide policymakers to very different solutions.

We use a difference-in-differences strategy in the context of San Francisco's rent control referendum in 1994 to find that increased exposure to expanded rent control decreased incidents of domestic violence severe enough to merit hospitalization. We examine trends in assaults on men and hospitalization patterns to rule out that these effects are due to changes in overall crime or hospital-going patterns. Furthermore, our analysis of demographic data from the U.S. Census suggests that moving between neighborhoods does not drive these results. Our estimated effects reflect changes within pre-existing relationships and changes in couple formation. They should be interpreted as aggregate effects of rent control, rather than changes in the propensity to be affected for an individual woman. These results are consistent with the financial strain model of domestic violence and align with the view of some policy makers that rent control policies will benefit the victims of domestic violence.

CHAPTER 3

Rational Eviction: How Landlords Use Evictions in Response to Rent Control

(with Eilidh Geddes)

3.1. Introduction

As housing prices rise, more cities are turning to rent control policies with the goal of ensuring long-term affordable housing. In a typical rent control policy, leases must be renewed at statutorily limited rent increases. Rent control policies reduce the returns from operating in the rental market, creating well-studied incentives to leave the rental market. Many rent control policies include "vacancy decontrol" provisions, which allow landlords to reset rents to market rates when tenants move. These policies limit the reductions in returns to operating in the rental market, but create incentives to induce tenant turnover, either through tenants moving or evictions. The more tenants move, the more often a landlord can raise rents to market rates.

In this paper, we examine whether a rent control expansion led to more eviction notices and increased complaints about wrongful evictions in San Francisco. We use a differences-in-differences research design exploiting ZIP code-level variation created by a 1994 ballot referendum that led to the removal of a rent control exemption for small buildings built before 1980. We find substantial increases in both the number of eviction notices and wrongful evictions claims filed with the San Francisco Rent Board in areas more affected by the policy change.

We first document a sharp increase in eviction notices reported to the Rent Board overall starting in 1995, when rent control is expanded. In San Francisco, under the Rent Ordinance, unless the tenant is in breach of the lease in some way, legal evictions require the landlord or a member of their immediate family to occupy the unit after the tenant leaves, remove all rental units from the rental market under the Ellis Act, or demolish the rental unit. The rise in eviction notices is concentrated in the types of evictions that most related to rent control-based incentives: owner move-in evictions and Ellis Act evictions. These evictions remove units from the rental stock, at least in the short term. We do not observe large increases in evictions that landlords cannot directly control, such as evictions for non-payment.

We also observe a sharp increase in claims of illegal evictions. Illegal evictions would include a misrepresentation of the owner's plans to the tenant to convince them to leave earlier, incorrect notice being given around an eviction, or a self-help eviction where the landlord either removes the tenant directly or interferes with the habitability of the unit. While there are many ways for a landlord to legally evict, these generally do not allow the landlord to receive the immediate benefit of resetting the rent to the market rate. For example, with an Ellis Act eviction, the owner must remove all units from the rental market. In contrast, illegal or "wrongful" evictions may produce this benefit for the landlord.

We conduct a ZIP code level difference-in-differences analysis where we compare ZIP codes with more and fewer units that newly become eligible for rent control based on their age and size. We find an 83% increase in eviction notices filed with the Rent Board and a 125% increase in the number of wrongful eviction claims for ZIP codes with the average level of new exposure to rent control. The increase in evictions occurs gradually over the five years following the policy change, likely due to market rents increasing over time. Our largest effects occur in years in which rent control would be particularly binding.

We additionally document heterogeneity with respect to median income in a given ZIP code, where effects on low income ZIP codes are 66%-200% higher than effects in high income ZIP codes. This heterogeneity is important in the setting of rent control where the goal of the policy is to prevent lower income tenants from being displaced from their homes when rents rise. Understanding the distributional consequences is important for policy makers.

Our paper contributes to the literature on rent control (Autor et al. (2014b), Early (2000), Glaeser and Luttmer (2003b), Olsen (1972b), Sims (2007)) by examining a new measure by which landlords respond to rent control regulations and documenting that this response leads to formal complaints filed by tenants. Much of the literature looking at supply side responses focuses on how supply to the rental market is affected. The most closely related paper, Diamond et al. (2019b), uses the same institutional setting and similar research design and finds that landlords reduced the rental supply by 15% by selling to owner-occupants and redeveloping buildings while rent control limited mobility by 20%.

We make several further contributions beyond this work. First, we examine the behavior of landlords who stay in the housing market and the actions that they take that limit the positive displacement effects of rent control. Second, we show how use of vacancy decontrol as a tool in the rental market coincides with periods in which rent control is particularly binding. Finally, our work highlights heterogeneity across landlords in the propensity to use these tools; the rise in eviction notices we find exceeds the increase if pre-policy landlord behavior continued. Our results highlight that landlords are able to use the evictions process to circumvent rent control policies and that the use of these evictions proceedings spikes in periods of rapid rent increases.

Other work focuses on the effects of price changes on eviction behavior. Asquith (2019a) and Asquith (2019b) examine how landlords in San Francisco respond to price increases under rent control using bus lines as an instrument for prices. These papers find evidence of an increase in owner move in evictions, consistent with our findings, but no evidence of increases in Ellis Act evictions. There were additionally increases in condo conversions. Pennington (2021) finds that evictions in rent controlled units fall after rents decrease due to an increase in housing supply.

Our contributions are complementary to these findings: we document the effects of an expansion of rent control, which may affect the market differently than price changes, specifically study "wrongful" eviction claims, and highlight that different types of landlords may behave differently. In particular, small landlords may have higher ability to leverage owner move in evictions.

The rest of this paper is organized as follows. Section 3.2 discusses the institutional details of rent control in San Francisco. Section 3.3 describes the data, and Section 3.4 describes our empirical design. Section 3.5 describes our results where we find an increase in wrongful eviction claims and eviction notices following the removal of a rent control exemption. Finally, Section 3.6 concludes.

3.2. Background of San Francisco's Rent Control Policies

In 1979, San Francisco passed its Rent Ordinance. Buildings built and occupied before 1979 above a certain size had rent increases capped at 60 percent of the regional CPI[1]. Under this legislation, rents will likely fall below market rents over time. In 1979, there was an exception for small (less than five units), owner-occupied buildings from rent control. This exemption was lifted by a ballot referendum in 1994[2], expanding the number of units subject to rent control within San Francisco by about 68% for the average ZIP code. As of 2015, roughly 60% of San Francisco's rental stock is rent controlled (40% of the overall housing stock), largely because a sizable amount of the rental stock was built before 1979 (San Francisco Planning Department 2018).

San Francisco's rent control policy includes "vacancy decontrol," which allows landlords to reset the rent to market rate when a tenant leaves the unit. These regulations also introduce "just cause" regulations, where landlords must have grounds for a lease termination or eviction[3]. Landlords can only evict tenants for one of 16 specific legal reasons, and they must have an "honest intent, without ulterior motive".

Landlords can then raise rents to market only in a limited number of situations. First, they can do so when a tenant leaves the apartment. For this reason, landlords might be incentivized to offer a cash buyout to existing tenants to induce them to move. Landlords can also perform an owner-move in eviction where they or an immediate family member

[1]The initial limit on rental increases was 7 percent per year which was lowered to 4 percent per year in 1984 and to 60 percent of CPI in 1992

[2]The ballot referendum, Proposition I, passed in a close election (51% vs 49%) on November 8, 1994, and went into effect on December 22, 1994. The rent charged on May 1, 1994, was considered to be the "base rent" for newly rent controlled buildings.

[3]Just cause regulations are generally passed at the same time as rent control policies. In San Francisco, the just cause regulations passed with the 1979 Rent Ordinance

begin occupying the unit[4]. On one hand, this is less financially beneficial, since while they can raise the rent after an owner move-in eviction, they would do so on themselves or an immediate family member, and if they re-let the unit before three years has passed, the rent control regulations prohibit them from raising the rent above what the prior tenant would have paid[5]. On the other hand, it is easy for landlords to evade these laws. The Rent Board only audits 10% of units each year, and if a unit's rent price is found to be higher than what is allowed based on owner tenancy and rent control laws, the landlord will be assessed a fee of up to $1,000 per month when excess rent was charged[6]. There additionally are potentially benefits associated with leasing to a new tenant; if the previous tenant was likely to reside in the unit for many years, it may be financially advantageous to find a new tenant.

Panel A Figure 3.2 shows the total number of evictions per year in San Francisco split by type. Prior to the policy change, all eviction types trend comparably. After the policy change, nuisance and non-pay evictions stay constant, while there are increases in owner move-in, breach, and "other" evictions. The evictions that rise are "no-fault" evictions, which landlords directly control and we would expect to rise most in response to a rent

[4]In order for an owner move-in eviction to be valid, the owner or relative must move into the unit within three months of an eviction notice, and must occupy the unit for at least 36 continuous months. An immediate family member moving into the unit can be grounds for an owner move-in eviction, but only if the owner already resides in the building. Further, if the owner moves out of the unit before their three-year required tenancy is up, the rent charged to the next tenant must not exceed what would have been charged to the tenant who lived in the unit prior to the landlord moving in. After December 18, 1998, an owner move-in eviction requires that the unit be designated as the owner's unit for any subsequent owner move in evictions

[5]https://sfrb.org/topic-no-204-evictions-based-owner-or-relative-move

[6]A punishment of up to six months of imprisonment in the County Jail is also possible for landlords who violate the rent control laws.

control policy. However, these eviction types are also subject to many regulations, and it is possible that many of these eviction notices were wrongfully handled by landlords.

Wrongfully handled "no-fault" evictions are one reason why a tenant might file a wrongful eviction claim with the Rent Board. Other reasons could include being forced out because of repair issues or because of landlord harassment. One of the first steps in fighting a wrongful eviction is filing a claim with the San Francisco Rent Board, which handles reports of evictions that violate Rent Ordinances. If a landlord is found to have wrongfully evicted a tenant in violation of the Rent Ordinances, they face financial penalties, which can be significant.

In general, the eviction process begins when a landlord serves an eviction notice to their tenant. The landlord has ten days to file the eviction notice with the Rent Board, but for no-fault evictions, more notice is required. For example, Ellis Act evictions require the eviction notice to be filed with the Rent Board 120 days before the withdrawal date. If the tenant does not move out before the withdrawal date, then the landlord can file and serve an "Unlawful Detainer Summons and Complaint" to the tenant in order to remove them from the unit, to which the tenant has five days to respond in court. The court will set a trial date and if the landlord wins, a sheriff will carry out the eviction.

A wrongful eviction claim can also be made at any time during the eviction process[7], though many grounds for wrongful eviction inherently require a notice to be served first [8] There is no financial cost to filing a wrongful eviction claim; however, the tenant must

[7]A tenant may also choose to sue the landlord for damages, though this would likely be a more expensive route and may require waiting until the eviction is complete, rather than fighting the eviction so that it never takes place.

[8]For example, a wrongful eviction claim may allege that the "just cause" written on the eviction notice is incorrect or does not apply. The eviction notice must have been served before this wrongful eviction claim can be submitted since it refers to the eviction notice itself.

be aware of the option to do so and pay the hassle cost of doing so. Once a wrongful eviction claim is filed, the Rent Board determines whether there is evidence of an unlawful eviction. If so, there will be an investigatory hearing before an Administrative Law Judge, who prepares a report for the Rent Board of Commissioners. The Rent Board of Commissioners will then determine whether to take further action (including making a referral to the District Attorney for criminal prosecution [9]). Wrongful eviction claims vary with tenant incentives. The longer a tenant stays in a rent controlled unit, the lower their relative rent, and the higher the benefit to filing a wrongful eviction claim if they are evicted.

3.3. Data

In this section, we describe the data that we use to measure the expansion of this rent control policy at the ZIP code level and the response to this policy by landlords in terms of eviction notices and wrongful eviction claims filed by tenants.

3.3.1. Measuring Rent Control

To understand the level of rent control treatment in each ZIP code, we use data on each unit's address, the number of units in the building, and the year the building was built for all residential units in the San Francisco Assessor's Secure Housing Roll from 1999. Since the rent control policy only affected small buildings with fewer than five units that were owner-occupied and built before 1980, we categorize buildings as newly rent-controlled if they have less than five units and were built in 1979 or earlier. The policy only affected

[9]The Rent Board does not publish statistics on the proportion of eviction notices or wrongful eviction claims that result in criminal prosecution or other legal action.

owner-occupied units, so we further restrict our definition to buildings with at least two units. We do not have an exact owner-occupancy measure in 1994 when the policy was passed. A map of our measure by ZIP code is available in Figure 3.1.

Figure 3.1. ZIP Code Level Exposure to Change in Rent Control Policy

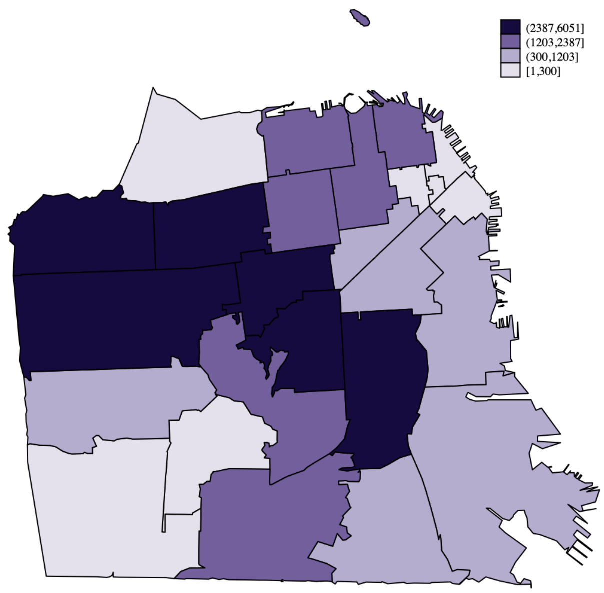

Notes: This map depicts exposure at the ZIP code level for the San Francisco policy change. Exposure is measured as the number of units in a ZIP code in buildings with 2-4 units that were built prior to 1979.
Data source: 1999 San Francisco Assessor's Secure Housing Roll and authors' calculations.

Some buildings may be counted as treated when they were actually rent-controlled both before and after the policy change due to not being owner-occupied. This miscategorization may attenuate our results[10]. In Appendix Section 2, we discuss an alternative measure of exposure that attempts to adjust for this owner occupancy element,

[10]This attenuation results from the fact that all units we classify as untreated were definitively not treated, but some units we classify as being treated may not have been treated in reality due to not being owner occupied. The treatment effect based on number of treated units may actually be higher on a per-unit basis because the number of treated units is overestimated.

defining an alternative treatment measure as the number of buildings with between two and four units, which were built before 1980, and which were owner occupied in 1999. Our results are robust to using this different measure. See Appendix Section 2 for a detailed discussion of the creation of the alternative treatment exposure measure, and Appendix Section 1 for a detailed discussion of the data cleaning procedure.

We construct further measures of treatment that address the discrepancies that may result from our treatment data coming several years after the implementation of the policy. First, we account for the fact that small rent controlled buildings may have been torn down and replaced by single family homes and assume all new single family homes build between 1995 and 1999 replaced a duplex. Second, we construct different measures to address condo conversions, including assuming all condos build before 1980 were converted in the post period, all condo modifications were condo conversions that replaced rent controlled units, and all new condos replaced rent controlled units. Finally, we construct a measure that assumes all new construction in the post period replaced rent controlled units. These measures are highly correlated.

3.3.2. Measuring Evictions

We measure the response of landlords to the rent control expansion using the ZIP code level number of eviction notices and wrongful eviction claims. Both measures originate with the San Francisco Rent Board. Our eviction notices data have information on the address and reason for the eviction. During our sample period, these data are incomplete in terms of the reason for the eviction and the ZIP code of many units; we discuss these limitations and our attempts to rectify the missing information in Appendix Section 3.

Both wrongful eviction claims and eviction notices can only be submitted by individuals who already live in or are landlords of rental units subjected to rent control. When rent control expands, even if the total number of evictions stays constant, the number of wrongful evictions claims could increase solely due to more tenants living in rent controlled apartments and therefore allowed to submit a wrongful eviction claim to the Rent Board. We discuss the implications of this and rule out that this mechanical effect drives results in Section 3.4.

Panels B, C, and D of Figure 3.2 show the average number of eviction notices, owner move in eviction notices, and wrongful evictions claims by ZIP code for ZIP codes in the lowest, middle, and highest tercile of the number of units newly treated by rent control as determined by the age and unit size of the building. Before the referendum's passage in 1994, all three terciles had roughly constant average reports of eviction notices and wrongful eviction claims. Starting in 1996, we see sharp increases in the number of wrongful eviction claims made in ZIP codes with medium or high levels of units newly exposed to rent control policies. However, in ZIP codes where relatively few units are newly exposed to rent control, we see that wrongful eviction claims remain roughly constant.

While reliable ZIP code level data on eviction notice types is not available before 1994, we plot owner move in evictions by treatment terciles starting in 1994 when the data start being available. From this, we see that the increase in owner move in evictions in panel A is driven largely by ZIP codes in the highest treatment tercile.

Summary statistics for our analysis dataset are displayed in Table 3.1. Panel A reports characteristics at the ZIP code level from the 1990 Census. We report these statistics for ZIP codes overall, for ZIP codes below the median level of exposure (the "low treatment"

Figure 3.2. Eviction Notices and Wrongful Eviction Claims Over Time

(a) Types of Eviction Notices

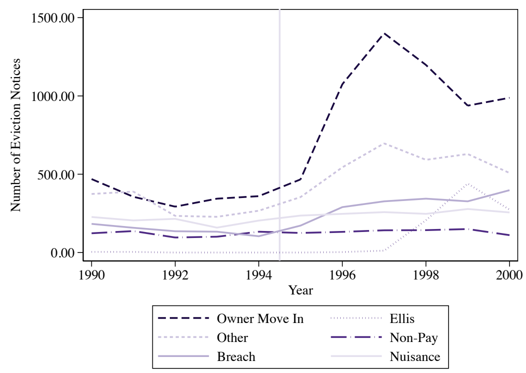

(b) Owner Move In Eviction Notices By Treatment Terciles

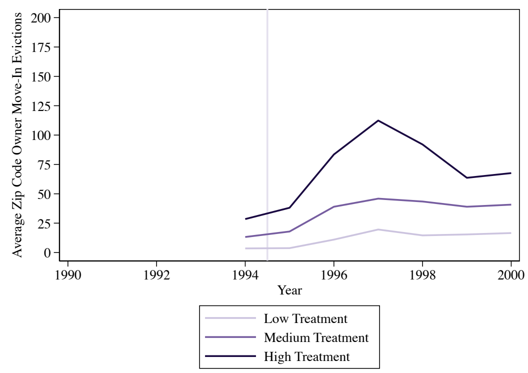

(c) Eviction Notices By Treatment Terciles

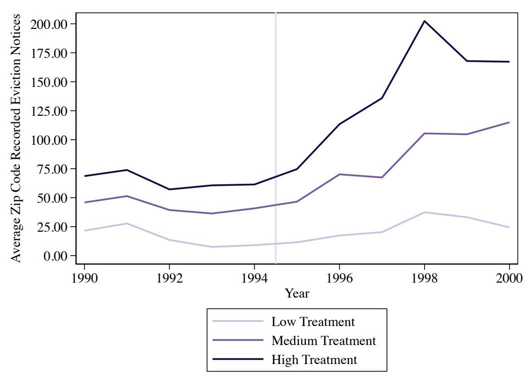

(d) Wrongful Eviction Claims By Treatment Terciles

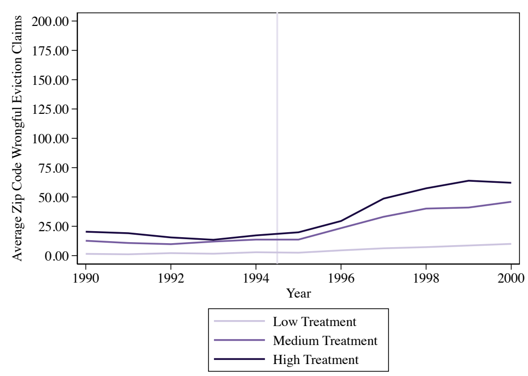

Notes: Panel A breaks down the number of eviction notices filed with the San Francisco by type over time. It is not possible to further breakdown evictions classified as "Other". Panel B splits owner move in eviction notices by treatment terciles for ZIP codes. These data are not available with any reliability at the ZIP code level before 1994. Panels C and D show the average number of eviction notices and the average number of alleged wrongful eviction reports made to the San Francisco Rent Board, respectively, by year for three treatment tercile groups. In Panels C-D, low treatment ZIP codes are those in the lowest tercile of units newly exposed to rent control after the policy change. Middle treatment ZIP codes are those in the middle tercile. High treatment are those in the highest tercile. Our treatment is measured by the number of units that are newly rent controlled. They are in buildings that were built before 1980 and have 2-4 units. The timing of the policy change is marked by a vertical line.

group), and for ZIP codes above the median level of exposure (the "high treatment" group). We also report the difference for each characteristic between the control and treatment groups and test whether this difference is statistically significant. We find no differences between our less treated and more treated ZIP codes.

Panel B examines housing characteristics. We see that there are differences across our highly treated and less treated ZIP codes in both pre- and post-period housing characteristics. Our highly treated ZIP codes both have more newly rent controlled units as well as more rent controlled units before the policy change. There are also differences in levels of the number of eviction notices and wrongful eviction claims before the policy. This difference should not be surprising, given that there is visual evidence of these differences in levels in Figure 3.2. We see growing differences in the number of eviction notices and wrongful eviction claims post policy.

Table 3.1. Comparison Between Low and High Treated ZIP Codes

Panel A: Demographics

Variable	(1) Full Sample	(2) Low Treatment	(3) High Treatment	(4) Difference
Median Income	45,698.680	44,922.082	46,415.539	1,493.455
	(14,783.986)	(19,287.252)	(9,745.483)	(6,182.045)
1990 Median Rent	869.280	869.000	869.538	0.538
	(218.606)	(296.193)	(123.116)	(91.947)
2000 Median Rent	1,006.360	1,028.833	985.615	-43.218
	(346.603)	(476.965)	(175.382)	(145.819)
Share Black	0.120	0.167	0.077	-0.091
	(0.137)	(0.170)	(0.080)	(0.054)
Share White	0.555	0.514	0.593	0.079
	(0.179)	(0.197)	(0.160)	(0.072)
Share Owner Occupied	0.316	0.327	0.307	-0.020
	(0.235)	(0.300)	(0.167)	(0.098)
Share Welfare	0.108	0.122	0.095	-0.028
	(0.077)	(0.095)	(0.058)	(0.032)
Panel B: Treatments and Outcomes				
Total Housing Units	12,622.280	6,927.417	17,879.076	10,951.660***
	(7,874.735)	(5,213.865)	(6,060.996)	(2,256.641)
# Prev. Rent Controlled	4,882.080	2,412.167	7,162.000	4,749.833**
	(5,024.150)	(2,729.110)	(5,650.232)	(1,755.680)
# Treated	1,688.480	358.417	2,916.231	2,557.814***
	(1,715.413)	(374.742)	(1,534.438)	(439.752)
Pre-Policy Eviction Notices	41.040	21.283	59.277	37.994***
	(30.852)	(21.822)	(26.769)	(9.739)
Post-Policy Eviction Notices	84.173	37.458	127.295	89.837***
	(66.350)	(31.180)	(60.965)	(19.172)
Pre-Policy Wrongful Claims	10.320	3.767	16.369	12.603***
	(9.690)	(4.750)	(9.193)	(2.898)
Post-Policy Wrongful Claims	28.927	11.875	44.667	32.792***
	(27.003)	(11.557)	(27.870)	(8.428)
Observations	25	12	13	25

Standard errors in parentheses
$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$

Notes: This table reports the averages of key variables. Column 1 reports averages for the full sample, while Columns 2 and 3 report averages for ZIP codes with below and above the median level of treatment, measured by the number of newly rent controlled units. The difference between the ZIP codes with higher exposure to the policy change and lower exposure to the policy change is reported in Column 4. Stars represent p-values of a t-test comparing the difference in Column 4 to zero.

3.4. Theoretical Context and Empirical Design

We model landlords as making rational decisions about whether to attempt to evict a tenant by weighing the expected benefit of evicting against the expected costs of pursuing an eviction. There are two sources of financial benefits for landlords: they are able to immediately re-set the rent to an uncontrolled rent price and they are able to rent the apartment to a new tenant who may have a higher propensity to move, allowing future re-sets to occur. Since renting an apartment generally becomes more expensive over time, incentives to wrongfully evict also increase gradually over time and will be higher in times where the market rents are rising faster than the allowed increase. Newly rent controlled units have additional increased incentives to turn over tenants, as the existing rent prices will not incorporate the value of future rent control.

Landlords also have the option to fully comply with regulations or willfully attempt a wrongful eviction. Wrongful eviction brings the highest benefit to the landlord, since legal evictions require renting to family members (owner move in evictions), buying out the lease, or removing all units from the rental market[11].

The costs of evicting include the paperwork costs of initiating the eviction, the likelihood of the tenant fighting the eviction and winning, and the cost to the landlord if that tenant successfully challenges the eviction[12]. A landlord is less likely to wrongfully evict a tenant who the landlord anticipates will fight the eviction or has the means to win a

[11]Legal evictions include evicting in order to sell all units in a building (Ellis Act Eviction), paying the tenant to leave voluntarily, or having the landlord or an immediate family member move into the unit. Landlords can also legally evict for breach of contract if tenants do not pay their rent. Apart from non-pay evictions, all legal evictions either present new costs or do not include the benefit of resetting to market rent.

[12]There are other costs to landlords of evicting, such as the cost of retaining an eviction specialist, which are constant across neighborhood and rent control status.

larger settlement in the ensuing legal fight. Tenants with lower incomes or who live in neighborhoods with lower median incomes may be even more likely to face a wrongful eviction, since their landlords may believe that lower income is correlated with a lower likelihood of hiring legal counsel and effectively navigating the legal system.

Altogether, we analyze three hypotheses. First, we expect that the rent control policy will lead to increased evictions. Next, we hypothesize that the increase in eviction notices and complaints will be larger in periods when landlords have more incentives to evict and that newly rent controlled units will evict tenants at higher rates than previously controlled units in the short run. As general equilibrium forces play out and some landlords leave the rental market, we expect the eviction level to stabilize, possibly returning back to its initial level. Finally, we explore heterogeneity with respect to median income of ZIP codes, to understand whether individuals who live in low income neighborhoods are even more likely to receive an eviction notice or submit a complaint of wrongful eviction.

3.4.1. Effect of Rent Control on Eviction Notices

To explore whether ZIP codes with more newly rent controlled units experience increased eviction rates, we use a continuous treatment difference-in-differences design. This design exploits variation across ZIP codes in the number of units who become exposed to rent control policies after the passage of the voter referendum in late 1994.

The parameter we identify is the average causal response (ACR), which indicates the additional increase in evictions due to an additional treated unit (Angrist and Imbens (1995)). In order to identify the ACR, we must make three assumptions: the Stable

Unit Treatment Value Assumption (SUTVA), parallel trends, and the homogeneity of the treatment effect across differently treated groups (Callaway et al. (2021)).

The Stable Unit Treatment Value Assumption requires that there are no spillovers between our treatment groups. This assumption could be violated in our setting if there are either anticipatory effects or spillovers across different ZIP codes. We find no evidence of anticipatory effects in our raw data plots in Figure 3.2. Additionally, given that the ballot referendum passed by a small margin, it is unlikely that tenants or landlords would have known whether the referendum would pass ahead of time (Harrison 1994).

A bigger concern is spillovers across units. In order to interpret our results as causal effects, we must assume units in less treated areas are not increasingly (or decreasingly) evicted due to price equilibrium effects caused by the policy change. This is more likely to be true in areas that were previously heavily rent controlled, since the high stock of rent controlled units would lead to stronger equilibrium effects. While we cannot measure overall rent effects due to lack of data, we can control for the number of previously treated units. We also can sign the size of the bias under the assumption that removing more units from market rents would create upward pricing pressure on market rents. This pricing pressure would create incentives for landlords in all rent controlled units to evict tenants, so if neighborhoods that are less treated by the policy still experience an increase in evictions due to the policy, our estimated effects will be a lower bound.

We also make a parallel trends assumption which requires that neighborhoods would have experienced similar trends in eviction levels had they been assigned a different number of treated units. Our raw data plots support this assumption.

Finally, we must assume the homogeneity of treatment effects across differently treated groups. Regardless of the number of rent controlled units of ZIP codes, the treatment effect of adding an additional rent controlled unit must be the same. This assumption would be violated if ZIP codes (or the individuals who live there) could select into different levels of treatment, since the covariates of individuals in neighborhoods who selected into different levels of treatment may affect their potential treatment effect. If this was the case, we would expect to see differences in characteristics potentially related to effect sizes based on the number of treated units. We find no evidence of differences in characteristics across groups in Panel A of Table 3.1, although we note that we are not powered to detect small differences. In Appendix Table A2, we examine these relationships in a linear regression framework and also find no evidence of demographic characteristics predicting treatment.

To identify the ACR, we estimate the following linear regression model:

$$Y_{it} = \alpha_i + \tau_t + \beta \cdot \text{Post}_t \times \text{Number Treated}_i + \epsilon_{it} \tag{3.1}$$

where α_i are ZIP code level fixed effects, τ_t are year fixed effects, Post_t is an indicator for whether the year is 1995 or after, and Number Treated_i is the number of newly treated units at the ZIP code level, calculated as discussed in Section 3.3.

We do not include time varying ZIP code controls due to lack of availability of data during the time period we study. We additionally estimate specifications that include treatment-tercile-specific linear time trends to bolster our parallel trends assumption and specifications that include the interaction between the number of previously treated units and the post period to account for spillovers.

3.4.2. Dynamics

To explore the hypotheses that the effects occur gradually over time, and possibly decline after the market has had time to react, we break down the "pre" and "post" periods into individual years in an event study specification. We interact each year with the number treated units in a given ZIP code. There is a single treatment time in our setting, so event time and calendar time are equivalent. Our event study specification is:

$$Y_{it} = \alpha_i + \tau_t + \sum_{t \neq 1994} \beta_t \cdot \text{Year}_t \times \text{Number Treated}_i + \epsilon_{it} \tag{3.2}$$

We normalize to 1994, the year that the policy passed.

3.4.3. Heterogeneity by Median Income

Finally, we explore heterogeneity by median income in the 1990 census. We group ZIP codes into whether their median income in the 1990s Census is above or below the median across ZIP codes. We estimate the following specification:

$$\begin{aligned} Y_{it} = \alpha_i + \tau_t + \gamma_t \times \text{Low Income}_i + \beta_1 \cdot \text{Post}_t \times \text{Number Treated}_i + \\ \beta_2 \text{Post}_t \times \text{Number Treated}_i \times \text{Low Income}_i + \epsilon_{it} \end{aligned} \tag{3.3}$$

where Low Income_i is an indicator for whether a ZIP code's median income was below the median for San Francisco in the 1990 Census and γ_t are year-high income fixed effects.

3.5. Results

3.5.1. Effect of Rent Control on Eviction Notices

Table 3.2 reports our coefficient estimates from Equation 3.1. We examine the effects of rent control expansion on eviction notices in Panel A and on wrongful eviction claims in Panel B. Column 1 reports the main difference-in-differences effect in thousands of treated units. For every additionally 1,000 newly rent controlled apartments, there are 20.07 additional eviction notices filed in that ZIP code and an additional 7.632 wrongful eviction claims. Since there were about 1,688 newly rent controlled units in each ZIP code on average, these effects constitute an 83% and 125% increase over the pre-treatment average level of evictions for the averagely-treated ZIP code[13].

In Column 2, we add treatment tercile specific fixed effects. In Column 3, we control for the effect of previously rent controlled units by including the interaction term between the number of previously rent controlled units and the post period. Our results are robust to the inclusion of these additional controls.

We cluster standard errors at the ZIP code level, our level of treatment variation. Since there are only 25 ZIP codes, we report p-values from a wild cluster bootstrap, which more accurately estimates clustered standard errors when the number of clusters is small (Cameron, Colin A. et al. (2008); Canay et al. (2021)). P-values for each relevant regression estimate are calculated using the routine provided by Roodman et al. (2019).

[13]In Appendix figure A9, we show that wrongful evictions increase as a constant fraction of eviction notices.

3.5.2. Dynamics

Figure 3.4 reports our event study coefficient estimates of Equation 3.2. While there is very little difference between the high and low exposure ZIP codes before the policy changed, supporting the parallel trends assumption, after the policy changed, we see that both eviction notices and wrongful eviction claims rise only in ZIP codes with higher levels of exposure. The effect increases over time, consistent with the changing incentives to evict. We report our difference-in-difference estimates in the dashed line.

Figure 3.4. Effect of Removing Rent Control Exemptions On Eviction Notices and Wrongful Eviction Claims

(a) Effect on Eviction Notices

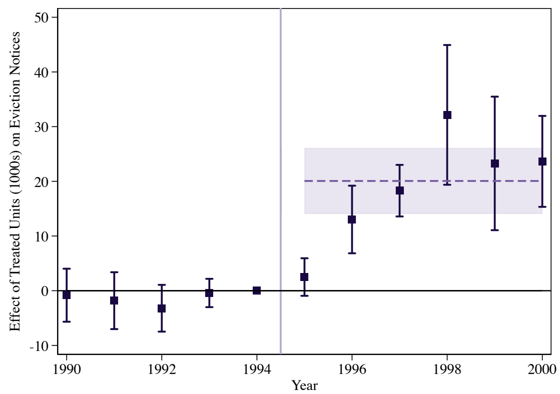

(b) Effect on Wrongful Evictions Claims

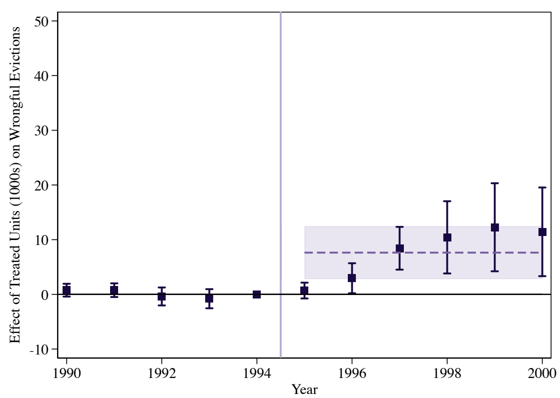

Notes: This figure shows our event study coefficient estimates on the interaction between year dummies and the number of treated units in a ZIP code. We normalize 1994 to be zero. Error bars shown are for the 95% confidence interval. The difference-in-differences estimate for the interaction of the post-period with the number of treated units is shown as a grey dashed line with the 95% confidence interval shown as a shaded region on the graph. Standard errors are clustered at the ZIP code level. Effects are scaled to be the effect per 1,000 treated units. The average number of treated units in a ZIP code is 2,791.

Evaluating the direct effects of the policy on rent is difficult as rent data at the ZIP code level from the 1990s is difficult to find. In Appendix Figure A8, we look at the effects of the policy on rent using ZIP code level rents from the 1980, 1990, and 2000 Census and find weak evidence that rents in ZIP codes that were heavily exposed to this rent control expansion experienced small declines relative to other ZIP codes, possibly due to the direct effects of the policy change.

We focus primarily on the short-run effects, as the parallel trends assumptions for our identification strategy become harder to justify in the longer run. However, in Figure 3.6, we present suggestive evidence on the longer term effects on eviction notices and compare these effects to the difference between aggregate rent prices and statutorily allowed rent increases. The effects converge to the expected mechanical increase of 57% by the early 2000s and are higher in periods where rent prices are rising faster than the statutorily allowed rent increases.

3.5.3. Heterogeneous Effects By Income

Table 3.2 Columns 4-6 report the results of Equation 3.3, evaluating whether low income ZIP codes are more likely to experience higher evictions. We find that low income ZIP codes are significantly more affected by the increases in rent control, both in terms of eviction notices and wrongful eviction claims. Low income neighborhoods experience 15 additional evictions per 1,000 treated units, compared to only 9 additional eviction notices in high income ZIP codes. Likewise, low income tenants file nearly 7 additional wrongful eviction complaints per 1,000 treated units, compared to only 2.5 in high income ZIP codes. Low income neighborhoods are experiencing eviction notices at rates at a level 67% higher than high income ZIP codes, and many evictions are handled incorrectly by

Figure 3.6. Effect of Removing Rent Control Exemptions On Eviction Notices and Allowed Rent Increases

(a) Allowed Rent Increase vs Rent CPI

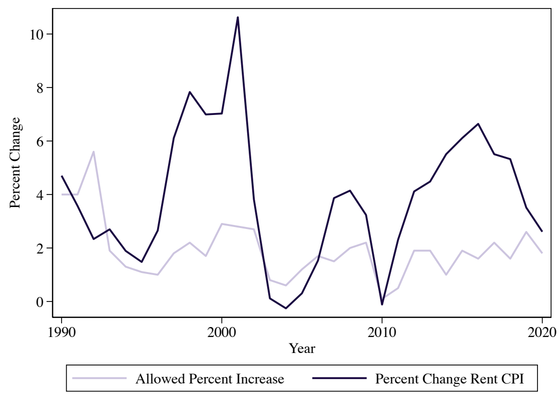

(b) Effect on Eviction Notices Over Time

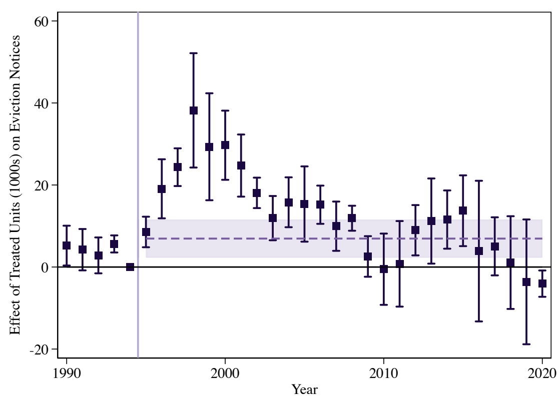

Notes: Panel A shows allowed rent increases and San Francisco rent price index, both in percent changes year over year. The allowable increase is set to be 60% of the percent increase in the consumer price index in San Francisco, calculated annually for the one-year period ending each October 31st. Panel B shows our event study coefficient estimates on the interaction between year dummies and the number of treated units in a ZIP code over the longer time horizon from 1990 until 2020. We normalize 1994 to be zero. Error bars shown are for the 95% confidence interval. The difference-in-differences estimate for the interaction of the post-period with the number of treated units is shown as a grey dashed line with the 95% confidence interval shown as a shaded region on the graph. Standard errors are clustered at the ZIP code level. Effects are scaled to be the effect per 1,000 treated units. The average number of treated units in a ZIP code is 1,688.

landlords. We additionally plot equivalent event study estimates in Figure 3.8 where we compare the effect over time for low and high income ZIP codes.

Figure 3.8. Heterogeneous Effects By Income

(a) Effect on Eviction Notices

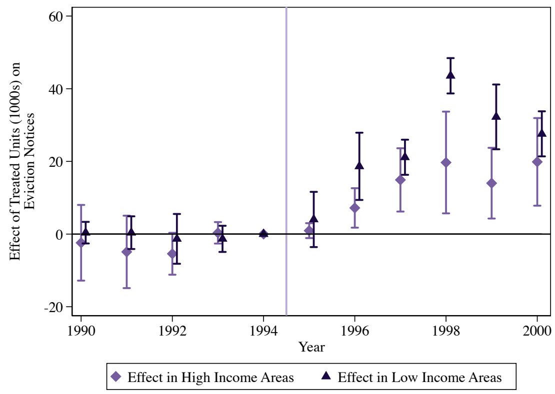

(b) Effect on Wrongful Evictions Claims

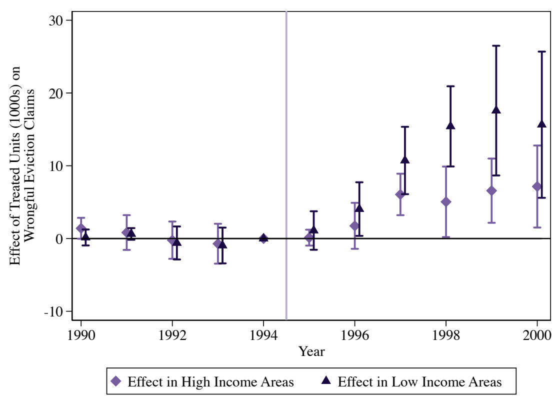

Notes: This figure shows our event study coefficient estimates on the interaction between year dummies and the number of treated units in a ZIP code (light purple) and this added to the interaction of these estimates with the triple interaction with low income (dark purple). We normalize 1994 to be zero. Error bars shown are for the 95% confidence interval. Standard errors are clustered at the ZIP code level. Effects are scaled to be the effect per 1,000 treated units. The average number of treated units in a ZIP code is 1,688.

Data Sources: 1999 San Francisco Assessor's Secure Housing Roll (Office of the Assessor-Recorder (1999)), San Francisco Rent Board Evictions Notices and Wrongful Evictions Claims (San Francisco Rent Board (2021))

3.5.4. Additional Robustness Checks

Beyond the alternative specifications already discussed, we perform additional robustness checks in the appendix. First, we use a binary measure of treatment, which discretizes exposure, to address concerns that rent control exposure should not enter the specification linearly. We compare neighborhoods with above-median numbers of newly rent controlled units to those with below-median exposure. We find large effects comparable to our main specifications.

Second, we use an alternative measure of the number of newly rent controlled units to address concerns that we have mis-measured treatment by not accounting for the fact that buildings needed to be owner-occupied to be eligible for the exemption to rent control prior to 1994. In Appendix Section 2, we detail the steps we take to account for this regulatory feature. We find that use of this alternative measure of policy exposure does not substantively change our results; while our estimates change in size, this change is proportional to the change in our treatment variable.

Third, we explore robustness to alternative measures of treatment that account for potential changes in assessor records that may have occurred in response to the policy change. Appendix Tables A5 and A6 report our estimates using different measures of treatment that assume different units whose classification may have changed between 1994 and 1999 were treated. Our estimates across specifications are similar relative to the number of treated units.

Finally, we investigate whether our effects can be attributed to the mechanical increase in the number of units whose evictions would be reported to the Rent Board. Only units covered by the Rent Ordinance will generate either eviction notices or wrongful eviction claims. We compare the average number of eviction notices and wrongful eviction claims to the number of units covered by the Rent Ordinance. In 1994, there were 1,068 evictions recorded with the Rent Board across San Francisco (360 owner move in and Ellis Act evictions) and 122,052 total rent controlled units. This represents a 0.8 percent eviction rate for rent controlled units. By 1997, there are 2,836 evictions (1,412 owner move in and Ellis Act Evictions) with 191,840 rent controlled units, a 1.5 percent eviction rate, an almost doubling of the eviction rate. In Appendix Figure A6, we show a time series of the number of eviction notices and wrongful eviction claims, which shows a jump up in the ratio of evictions to rent controlled units following the policy change.

3.5.5. Discussion

Several behavioral changes could drive our results: landlords could increase lawful evictions, landlords could increase wrongful evictions, and tenants could change how they challenge evictions. In order to see our results, landlord behavior must be changing, since we see large increases in eviction notices, which are unlikely to be influenced by tenant behavior. These eviction notices are particularly concentrated in owner-move in evictions, which are in the direct control of the landlord.

This change in evictions could be driven either by an increase in wrongful or legal evictions. We don't find evidence of large changes in the proportion of wrongful eviction claims per eviction notice at the ZIP code level (there were 0.25 claims per notice pre-policy and 0.34 claims per notice post-policy), suggesting that landlords post-policy are not increasing the proportion of evictions that cause tenants to complain, although the absolute number of wrongful eviction claims increases.

Newly rent-controlled tenants may behave differently in response to an eviction notice than previously rent controlled tenants. This change could go in either direction; if the policy change coincided with greater awareness of tenant's rights, there could be an increase in the number of wrongful eviction claims for rent controlled tenants. However, newly rent controlled tenants may either have less incentive to file a claim (their rent is not as far removed from market rent as a long-term rent controlled tenant) or may be less aware of the regulations around just-cause evictions, making them less likely to file a wrongful eviction claim.

Table 3.2. Difference-in-Difference Estimates of the Effects of Rent Control on Wrongful Evictions Claims

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Eviction Notices						
Num Treated × Post	20.07***	15.39***	19.50***	15.24***	6.986	15.26***
	(2.926)	(3.776)	(3.044)	(2.374)	(4.214)	(2.323)
Num Treated × Post × Low Income				9.662***	11.99***	9.669***
				(2.651)	(3.637)	(2.677)
Treatment Tercile		X			X	
Prev. Treated Control			X			X
N	275	275	275	275	275	275
R^2	0.882	0.890	0.883	0.892	0.901	0.892
P-value Num Treated × Post	0	0.00400	0	0.00500	0.184	0.00500
P-value # Treated× Post ×Low Income	.	.	.	0.0160	0.0170	0.0120
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: Wrongful Eviction Claims						
Num Treated × Post	7.632***	7.090**	7.197***	4.191***	2.017	4.030***
	(2.344)	(2.761)	(2.506)	(1.184)	(2.364)	(1.152)
Num Treated × Post × Low Income				6.718**	7.403**	6.653**
				(2.705)	(3.300)	(2.834)
Treatment Tercile		X			X	
Prev. Treated Control			X			X
N	275	275	275	275	275	275
R^2	0.830	0.840	0.831	0.847	0.858	0.847
P-value Num Treated× Post	0	0.0690	0.00100	0.0170	0.414	0.0160
P-value # Treated× Post ×Low Income	.	.	.	0.167	0.177	0.223

Notes: This table shows estimates from difference-in-difference regressions. In Panel A, the outcome variable is wrongful eviction claims, and in Panel B, the outcome is eviction notices. For both panels, Column 1 shows the results from our preferred specification, a difference-in-differences regression using the number of treated units at the ZIP code level as a continuous treatment measure. Column 2 adds heterogeneity by median income. Columns 3 and 4 add treatment tercile specific linear time trends to the regression, and Columns 5 and 6 add 1990 median rent linear time trends. Treatment is measured in 1,000s of newly rent controlled units. Data Sources: 1999 San Francisco Assessor's Secure Housing Roll (Office of the Assessor-Recorder (1999), San Francisco Rent Board Evictions Notices and Wrongful Evictions Claims (Board (2005)), U.S. Census Bureau and Social Explorer (2022)

3.6. Conclusion

San Francisco's expansion of rent control in 1994 led to a dramatic increase in rent control. Landlords of newly rent controlled units faced new incentives to turn over the units in order to raise the rents to market levels. One mechanism through which they could do this was eviction, either lawfully or unlawfully. We study the effects of the policy change on both eviction notices and wrongful eviction claims and find substantial increases in both in ZIP codes that were heavily affected by the policy change.

San Francisco's rent control policy change lead to an increase of roughly 20 eviction notices and 7 wrongful eviction claims per 1,000 buildings treated with rent control. Since San Francisco has roughly 1,688 treated units per ZIP code and 25 ZIP codes, we estimate that the rent control policy lead to 840 more eviction notices and 320 more wrongful eviction claims than would have happened otherwise. These effects are fairly persistent; it takes about seven years for our effect sizes to fall to the levels that would be predicted by the increase in rent controlled units that are subject to reporting requirements to the San Francisco Rent Board. The effects are also concentrated in low-income areas, with effect sizes 66%-200% larger in those areas.

We are cautious about interpreting our results too broadly, as the policy change we study only affects a particular kind of rental unit: small buildings that are owner occupied. "Mom-and-pop" landlords may be more willing to get out of the rental business altogether in response to changes in rent control regulation. In particular, they may be positioned to take advantage of owner-move-in evictions in a way larger landlords are not. Additionally, small landlords may be more likely to skirt the legal restrictions, either because they are not aware of the correct legal proceedings or because they are more willing to take legal risk than a large scale landlords.

Even so, our estimates imply that San Francisco's rent control expansion is responsible for a large portion of the increase in evictions in the 1990s. The 840 additional evictions comprise 51% of the increase in evictions overall in San Francisco from 1994 to 2000.

References

Aizer, Anna, "The gender wage gap and domestic violence," *American Economic Review*, 2010, *100* (4), 1847–1859.

Alexander, Diane, Janet Currie, and Molly Schnell, "Check up before you check out: Retail clinics and emergency room use," *Journal of Public Economics*, 2019, *178*, 104050.

Allende, Claudia, Francisco Gallego, Christopher Neilson et al., "Approximating the equilibrium effects of informed school choice," Technical Report 2019.

Anderson, Michael and Jeremy Magruder, "Learning from the crowd: Regression discontinuity estimates of the effects of an online review database," *The Economic Journal*, 2012, *122* (563), 957–989.

Angelucci, Manuela, "Love on the rocks: Domestic violence and alcohol abuse in rural Mexico," *The BE Journal of Economic Analysis & Policy*, 2008, *8* (1).

Angrist, Joshua D. and Guido W. Imbens, "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association*, June 1995, *90* (430), 431–442.

Armstrong, Carrie, "JNC8 guidelines for the management of hypertension in adults," *American family physician*, 2014, *90* (7), 503–504.

Asquith, Brian, "Do Rent Increases Reduce the Housing Supply Under Rent Control? Evidence from Evictions in San Francisco," SSRN Scholarly Paper ID 3165599, Social Science Research Network, Rochester, NY August 2019.

Asquith, Brian J., "Housing Supply Dynamics under Rent Control: What Can Evictions Tell Us?," *AEA Papers and Proceedings*, May 2019, *109*, 393–396.

Autor, David H, Christopher J Palmer, and Parag A Pathak, "Housing market spillovers: Evidence from the end of rent control in Cambridge, Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Autor, David H., Christopher J. Palmer, and Parag A. Pathak, "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, June 2014, *122* (3), 661–717.

Autor, David H, Christopher J Palmer, and Parag A Pathak, "Ending rent control reduced crime in Cambridge," in "AEA Papers and Proceedings," Vol. 109 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2019, pp. 381–384.

Beaulieu, Nancy Dean, "Quality information and consumer health plan choices," *Journal of health economics*, 2002, *21* (1), 43–63.

Bellemare, Marc F and Casey J Wichman, "Elasticities and the inverse hyperbolic sine transformation," *Oxford Bulletin of Economics and Statistics*, 2020, *82* (1), 50–61.

Board, San Francisco Rent, "Rent Board Memorandum," 2005.

Bobonis, Gustavo J, Melissa González-Brenes, and Roberto Castro, "Public transfers and domestic violence: The roles of private information and spousal control," *American Economic Journal: Economic Policy*, 2013, *5* (1), 179–205.

Brassiolo, Pablo, "Domestic violence and divorce law: When divorce threats become credible," *Journal of Labor Economics*, 2016, *34* (2), 443–477.

Bundorf, M Kate, Natalie Chun, Gopi Shah Goda, and Daniel P Kessler, "Do markets respond to quality information? The case of fertility clinics," *Journal of health economics*, 2009, *28* (3), 718–727.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

———, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," *arXiv:2107.02637 [econ]*, July 2021. arXiv: 2107.02637.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik, "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Calvi, Rossella and Ajinkya Keskar, "'Til Dowry Do Us Part: Bargaining and Violence in Indian Families," 2021.

Cameron, Colin A., Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-based improvements for inference with clustered errors," *The review of economics and statistics*, 2008, *90* (3), 414–427.

Canay, Ivan A., Andres Santos, and Azeem M. Shaikh, "The Wild Bootstrap with a "Small" Number of "Large" Clusters," *The review of economics and statistics*, 2021, *103* (2), 346–363.

Card, David and Gordon B Dahl, "Family violence and football: The effect of unexpected emotional cues on violent behavior," *The quarterly journal of economics*, 2011, *126* (1), 103–143.

Cesur, Resul and Joseph J Sabia, "When war comes home: The effect of combat service on domestic violence," *Review of Economics and Statistics*, 2016, *98* (2), 209–225.

Chan, David C, Matthew Gentzkow, and Chuan Yu, "Selection with variation in diagnostic skill: Evidence from radiologists," *The Quarterly Journal of Economics*, 2022, *137* (2), 729–783.

Chartock, Benjamin L, "Quality Disclosure, Demand, and Congestion: Evidence from Physician Ratings," 2021.

Currie, Janet and Jonathan Zhang, "Doing more with less: Predicting primary care provider effectiveness," *The Review of Economics and Statistics*, 2021, pp. 1–45.

______ and W Bentley MacLeod, "Diagnosing expertise: Human capital, decision making, and performance among physicians," *Journal of labor economics*, 2017, *35* (1), 1–43.

______, ______, and Jessica Van Parys, "Provider practice style and patient health outcomes: The case of heart attacks," *Journal of health economics*, 2016, *47*, 64–80.

Dahlstrand, Amanda, "Defying Distance? The Provision of Services in the Digital Age."," 2021.

Diamond, Rebecca, Tim McQuade, and Franklin Qian, "The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–3394.

______, ______, and ______, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, September 2019, *109* (9), 3365–3394.

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

———, **Daniel Kessler, Mark McClellan, and Mark Satterthwaite**, "Is more information better? The effects of "report cards" on health care providers," *Journal of political Economy*, 2003, *111* (3), 555–588.

Early, Dirk W., "Rent Control, Rental Housing Supply, and the Distribution of Tenant Benefits," *Journal of Urban Economics*, September 2000, *48* (2), 185–204.

Fadlon, Itzik and Jessica Van Parys, "Primary care physician practice styles and patient care: Evidence from physician exits in Medicare," *Journal of health economics*, 2020, *71*, 102304.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Sources of geographic variation in health care: Evidence from patient migration," *The quarterly journal of economics*, 2016, *131* (4), 1681–1726.

Fletcher, Jason M, Leora I Horwitz, and Elizabeth Bradley, "Estimating the value added of attending physicians on patient outcomes," Technical Report, National Bureau of Economic Research 2014.

Gelles, Richard J, "Abused wives: Why do they stay," *Journal of Marriage and the Family*, 1976, *38* (4), 659–668.

Glaeser, Edward L and Erzo F P Luttmer, "The misallocation of housing under rent control," *American Economic Review*, 2003, *93* (4), 1027–1046.

Glaeser, Edward L. and Erzo F. P. Luttmer, "The Misallocation of Housing Under Rent Control," *American Economic Review*, September 2003, *93* (4), 1027–1046.

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

Gowrisankaran, Gautam, Keith Joiner, and Pierre Thomas Léger, "Physician practice style and healthcare costs: Evidence from emergency departments," *Management Science*, 2022.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano, "Do fiscal rules matter?," *American Economic Journal: Applied Economics*, 2016, pp. 1–30.

Gruber, Jonathan and Robin McKnight, "Controlling health care costs through limited network insurance plans: Evidence from Massachusetts state employees," *American*

Economic Journal: Economic Policy, 2016, *8* (2), 219–250.

Grytten, Jostein and Rune Sørensen, "Practice variation and physician-specific effects," *Journal of health economics*, 2003, *22* (3), 403–418.

Harrison, Jim, "SF Voters Pass Bond Measure For Art Museum - Approval of $42 million to refurbish old library," *The San Francisco Chronicle B5*, 1994.

Hermida, Ramón C, Diana E Ayala, Artemio Mojón, and José R Fernández, "Bedtime dosing of antihypertensive medications reduces cardiovascular risk in CKD," *Journal of the American Society of Nephrology*, 2011, *22* (12), 2313–2321.

Hidrobo, Melissa, Amber Peterman, and Lori Heise, "The effect of cash, vouchers, and food transfers on intimate partner violence: evidence from a randomized experiment in Northern Ecuador," *American Economic Journal: Applied Economics*, 2016, *8* (3), 284–303.

Holz, Justin E, Roman G Rivera, and Bocar A Ba, "Peer effects in police use of force," *American Economic Journal: Economic Policy*, 2020.

Hull, Peter, "Estimating treatment effects in mover designs," *arXiv preprint arXiv:1804.06721*, 2018.

Jr, Joseph J Doyle, Steven M Ewer, and Todd H Wagner, "Returns to physician human capital: Evidence from patients randomized to physician teams," *Journal of health economics*, 2010, *29* (6), 866–882.

Kenney, Catherine T and Sara S McLanahan, "Why are cohabiting relationships more violent than marriages?," *Demography*, 2006, *43*, 127–140.

Kwok, Jennifer, "How do primary care physicians influence healthcare? Evidence on practice styles and switching costs from Medicare," *Evidence on Practice Styles and Switching Costs from Medicare (July 22, 2019)*, 2019.

Li, Cairu, Gunnar Engstrom, Bo Hedblad, Goran Berglund, and Lars Janzon, "Blood pressure control and risk of stroke: a population-based prospective cohort study," *Stroke*, 2005, *36* (4), 725–730.

Manning, Wendy D, Monica A Longmore, and Peggy C Giordano, "Cohabitation and intimate partner violence during emerging adulthood: High constraints and low commitment," *Journal of Family Issues*, 2018, *39* (4), 1030–1055.

Mayo, Jennifer, "Navigating the notches: charity responses to ratings," in "Working

Paper" 2021.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of econometrics*, 2008, *142* (2), 698–714.

Miller, Sarah, "The effect of insurance on emergency room visits: an analysis of the 2006 Massachusetts health reform," *Journal of public Economics*, 2012, *96* (11-12), 893–908.

Molitor, David, "The evolution of physician practice styles: evidence from cardiologist migration," *American Economic Journal: Economic Policy*, 2018, *10* (1), 326–356.

Munyo, Ignacio and Martín Antonio Rossi, "The effects of real exchange rate fluctuations on the gender wage gap and domestic violence in Uruguay," Technical Report, IDB Working Paper Series 2015.

National Health Interview Survey 2011 Data Release

***National Health Interview Survey 2011 Data Release**, 2011.*

Office of the Assessor-Recorder, *"Secured Assessment Roll," Technical Report 1999.*

Olsen, Edgar O, *"An econometric analysis of rent control," Journal of Political Economy, 1972, 80 (6), 1081–1100.*

Olsen, Edgar O., *"An Econometric Analysis of Rent Control," Journal of Political Economy, November 1972, 80 (6), 1081–1100.*

Parys, Jessica Van, *"Variation in physician practice styles within and across emergency departments," PloS one, 2016, 11 (8), e0159882.*

Penn, Madeline, Saurabha Bhatnagar, SreyRam Kuy, Steven Lieberman, Shereef Elnahal, Carolyn Clancy, and David Shulkin, *"Comparison of wait times for new patients between the private sector and United States Department of Veterans Affairs Medical Centers," JAMA Network open, 2019, 2 (1), e187096–e187096.*

Pennington, Kate, *"Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco," Working Paper, 2021.*

Rennison, Callie Marie, *Intimate partner violence, 1993-2001, US Department of Justice, Office of Justice Programs, Bureau of Justice . . . , 2003.*

Rhoades, Galena K, Scott M Stanley, Gretchen Kelmer, and Howard J Markman, *"Physical aggression in unmarried relationships: the roles of commitment and constraints.," Journal of Family Psychology, 2010, 24 (6), 678.*

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb, *"Fast and wild: Bootstrap inference in Stata using boottest," The Stata Journal: Promoting communications on statistics and Stata, March 2019, 19 (1), 4–60.*

Sabety, Adrienne, *"The Value of Relationships in Healthcare," Available at SSRN 4191234, 2021.*

San Francisco Rent Board, *"Eviction Notices," DataSF, 2021.*

Schnell, Molly, *"Physician behavior in the presence of a secondary market: The case of prescription opioids," Princeton University Department of Economics Working Paper, 2017, 5.*

——— **and Janet Currie**, *"Addressing the opioid epidemic: is there a role for physician education?," American journal of health economics, 2018, 4 (3), 383–410.*

Simeonova, Emilia, Niels Skipper, and Peter Rønø Thingholm, *"Physician health management skills and patient outcomes," Journal of Human Resources, 2022, pp. 0420–10833R1.*

Sims, David P., *"Out of control: What can we learn from the end of Massachusetts rent control?," Journal of Urban Economics, January 2007, 61 (1), 129–151.*

Stanley, Scott M and Howard J Markman, *"Assessing commitment in personal relationships," Journal of Marriage and the Family, 1992, pp. 595–608.*

Stevenson, Betsey and Justin Wolfers, *"Bargaining in the shadow of the law: Divorce laws and family distress," The Quarterly Journal of Economics, 2006, 121 (1), 267–288.*

Sun, Liyang and Sarah Abraham, *"Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," Journal of Econometrics, 2021, 225 (2), 175–199.*

Svarer, Michael, Michael Rosholm, and Jakob Roland Munch, *"Rent control and unemployment duration," Journal of Public Economics, 2005, 89 (11-12), 2165–2181.*

Truman, Jennifer L and Rachel E Morgan, *"Nonfatal domestic violence," Washington, DC: US Department of Justice, Bureau of Justice Statistics, 2014.*

U.S. Census Bureau and Social Explorer, *"1990 US Census," Technical Report, Prepared by Social Explorer 2022.*

Vatter, Benjamin, *"Quality disclosure and regulation: Scoring design in medicare advantage," Technical Report, Working Paper 2021.*

Wedig, Gerard J and Ming Tai-Seale, *"The effect of report cards on consumer choice in the health insurance market," Journal of health economics, 2002, 21 (6), 1031–1048.*

Yoon, Tae Jung, *"Quality information disclosure and patient reallocation in the healthcare industry: Evidence from cardiac surgery report cards," Marketing Science, 2020, 39 (3), 636–662.*

APPENDIX A

Chapter 1 Appendix

A.1. Difference between impacts of achieving "higher quality" versus "higher quality, lower cost" status

The main results of this study show that patients respond more to the higher status ("higher quality, lower cost") versus the middle "higher quality" (but not lower cost) status. This section discusses possible reasons for the differences.

First, one may wonder whether patients correctly understand that the difference between "higher quality, lower cost" status and "higher quality" status is based on cost. To determine whether this is the case, I examine the impacts of status separately for patients who paid nothing out of pocket for their visit. If these patients know ahead of time that they stand to pay nothing for their visit, they should ignore the cost part of the status, since the information about cost provided is irrelevant. If "higher quality, lower cost" physicians are still prioritized over "higher quality" (not lower cost) physicians, it is either because patients misunderstand the cost portion of the status, or because the response is driven by page rank rather than responses to the information provided.

Appendix Figure A.1 displays the results of this exercise. Panel A compares the responses of patients who paid nothing for their visit to those who paid for 100% of their visit out of pocket for lower cost physicians who achieve "higher quality, lower cost" status upon passing the threshold and "lower quality" status if not. My findings show clear effects for both types of patients. Panel B compares the same for "not lower cost" physicians and shows no effect. These results suggest that either patients misinterpret the status information provided (perhaps thinking that "higher quality, lower cost" physicians are even higher quality than"higher quality" physicians) or that the impact of page rank drives the results.

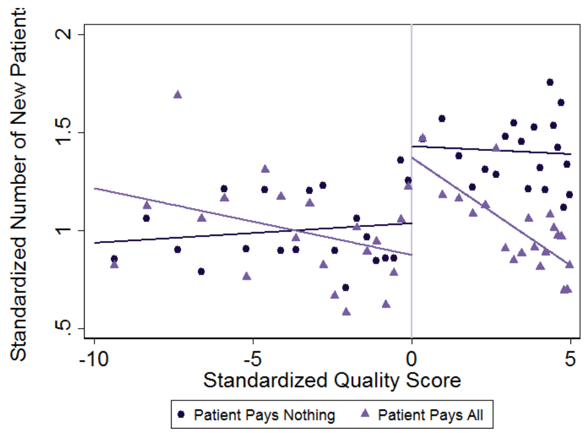

(a) Lower Cost Physicians

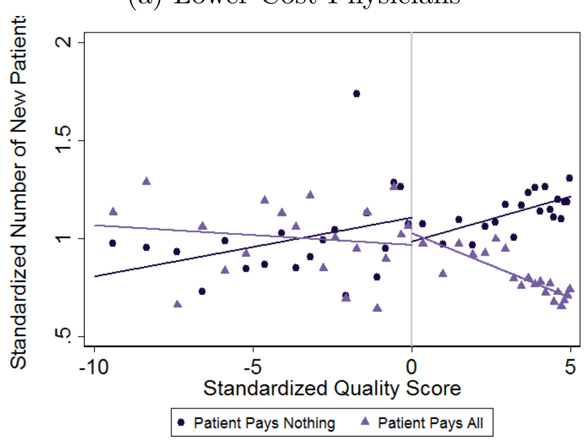

(b) Not Lower Cost Physicians

Figure A.1. Effects by Proportion Paid by the Patient
Notes: This figure displays impacts of physician status broken down by the amount patients paid, for patients who paid nothing out-of-pocket (Patient Pays Nothing) compared to patients who covered 100% of their care (Patient Pays All).

Second, one may hypothesize that the fact that regression discontinuity estimates are local drives the differences between effects. Perhaps patients prefer "higher quality" (not low cost) physicians over "higher quality, lower cost" physicians, but "higher quality" physicians have a lower capacity to take new patients. This would make treatment effects mechanically lower. Appendix Figure A.2 shows that providers who are not at

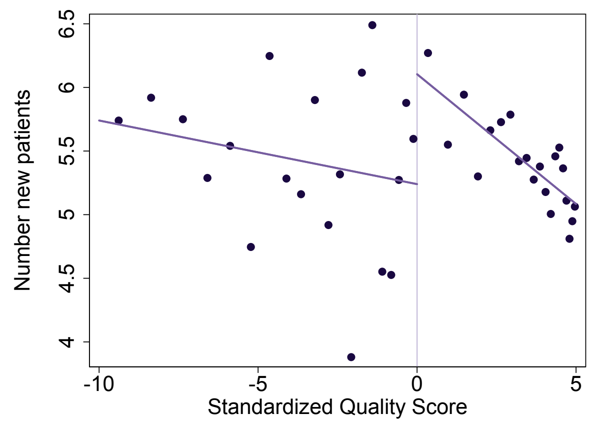

(a) Not at Capacity

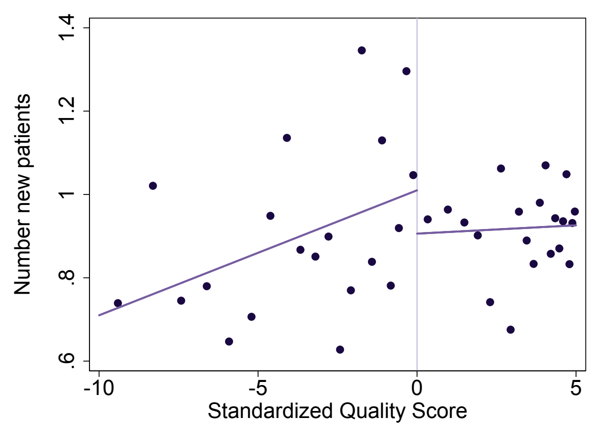

(b) At Capacity

Figure A.2. Effects by Provider Capacity
Notes: This figure displays impacts of the "higher quality, lower cost" status separately for physicians who are not at capacity in Panel A versus those who are at capacity in Panel B. Capacity is proxied as follows: A physician is considered to be at capacity if they did not see any new patients but saw a positive number of returning patients during the time period just before new statuses were released.

capacity (as proxied by the provider taking no new patients but having seen a positive number of returning patients during the pre-update period) experience large treatment effects, whereas physicians who are at capacity experience none. However, there were no significant differences in the proportion of physicians who were at capacity between these physician status groups. The impact difference between "lower cost" and not "lower cost" providers remains when breaking down to physicians who are not at capacity (see Appendix Figure A.2). Capacity constraints, therefore, do not explain the results.

A.2. Treatment Time Heterogeneity

Recent advances in the econometrics literature have pointed to homogeneity assumptions that can prove critical in two-way fixed effects designs such as the switcher design. Goodman-Bacon (2021) shows that the estimate of interest in a two-way fixed effect design estimates a weighted combination of treatment effects relative to various controls, where weights may have different signs. To overcome this estimation problem, researchers propose estimators that explicitly assign control groups to avoid bias caused by the aggregation implicit in OLS estimates of these two-way fixed effects designs (Callaway et al., 2021) (Sun and Abraham, 2021).

Novel estimators require a very specific setup, where individuals are compared to each other pre- versus post-treatment, so the specification in Equation (9) cannot be used to estimate impacts that are robust to treatment time heterogeneity because that specification requires the post-treatment indicator to interact with the change in quality and cost scores.

In lieu of estimating robust to treatment time heterogeneity (Equation (9)), which is not possible given currently available estimators, I use Callaway et al. (2021)'s method to estimate the following specification:

$$Y_{it} = \iota_i + \tau_i + \beta_1 1\{PostSwitch\}_{it} + \nu_{it} \tag{A.1}$$

The regression in Equation A.1 estimates the impact of switching physicians, regardless of the direction of the switch. To estimate impacts that are more comparable to the results of estimating the regression in Equation (9), I create subsets of groups of patients based on the direction of their switch. I first make a subset of patients who switched to a physician with a higher quality score, regardless of cost score. Second, I make a subset of patients who switched to a physician with a lower cost score, regardless of quality. Next, I create a subset based on both quality and cost, breaking patients down into four groups: those whose switch resulted in increased quality and cost scores, those whose switch resulted in increased quality scores but decreased cost scores, those whose switch resulted in increased cost scores but decreased quality scores, and those whose switch resulted in decreased quality and cost scores.

The regression in Equation A.1 can be estimated using both OLS and the Callaway and Sant'Anna (CS) estimators. Comparing the two estimation methods can provide information on whether heterogeneous effects based on treatment time are biasing the OLS results. Breaking down patients into different categories helps to evaluate whether, qualitatively, patterns match those discussed above.

Appendix Figure A.3 displays the results. The darker purple bars show the estimates of OLS regressions, while the lighter purple bars display the results of the analogous CS regression. Across the board, the CS estimates are larger than OLS estimates, suggesting that in this setting, OLS may bias impacts toward zero. The second point to notice is that the impact of switching a physician is always positive. That is, switching physicians results in higher spending, regardless of the cost or quality of the physician one is switching to. This result is also observed in the above estimation of Equation A.1; however, the

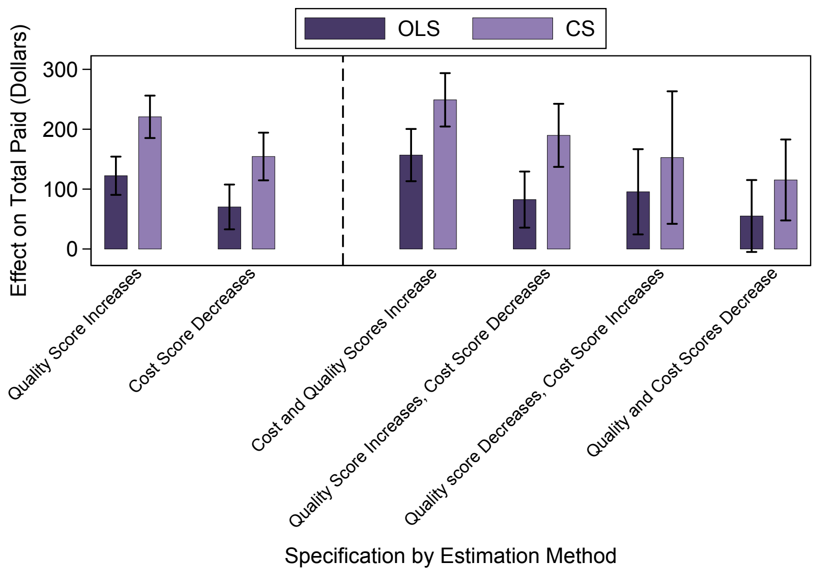

Figure A.3. Comparison of Effect of Switching within Sub-Groups: OLS vs CS
Notes: This figure plots impacts of switching over sub-groups, exploring the impacts of switching separately for patients whose switch resulted in them seeing a physician with a higher quality score, then for patients whose switch resulted in seeing a physician with a lower cost score. The next four sets of estimates subset further, first showing the impact of switching for those whose switch lead to an increase in quality and cost scores, then those whose switch lead to an increase in quality and a decrease in cost scores, third for those whose switch lead to a decrease in quality and an increase in cost scores, and finally for those whose switch lead to decreased quality and cost scores. While Callaway-Sant'Anna (CS) estimates are generally larger than OLS estimates, both sets of results point to the same takeaway, that switching to a physician with higher quality and cost scores results in higher spending than switching to a physician with lower quality and cost scores.

positive impact of switching on spending is captured by the relative year-fixed effects in x_it. Appendix Table A.1 displays the full regression results for the specifications with relative year-fixed effects and with post-switch indicators separately.

Qualitatively, the OLS and CS impacts look quite similar. First, consider a patient who switches to a new physician whose quality and cost scores are both higher. If the results in Section 9.1 are true, then this switch should result in the largest increase in

	(1) Total Paid	(2) Total Paid	(3) Total Paid	(4) Total Paid
$\Delta Q\times$ Post	1.646***	1.767***	1.692***	1.788***
	(0.553)	(0.555)	(0.550)	(0.552)
$\Delta C\times$ Post	-11.29***	-11.33***	-11.39***	-11.27***
	(3.707)	(3.701)	(3.694)	(3.689)
relative_switch_year==-4	-18.40		-20.08**	
	(19.31)		(8.939)	
relative_switch_year==-3	-14.64		-31.42***	
	(32.35)		(10.22)	
relative_switch_year==-2	-7.783		-36.64***	
	(46.40)		(11.92)	
relative_switch_year==-1	-88.66		-137.0***	
	(59.92)		(11.90)	
relative_switch_year==0	40.53		-24.09*	
	(73.44)		(14.06)	
relative_switch_year==1	36.41		-40.06**	
	(89.51)		(19.48)	
relative_switch_year==2	42.54		-49.45**	
	(102.3)		(19.67)	
relative_switch_year==3	63.01		-41.50*	
	(117.0)		(22.73)	
relative_switch_year==4	91.08		-35.19	
	(133.1)		(31.20)	
relative_switch_year==5	151.3		7.701	
	(147.2)		(32.64)	
post_switch		96.50***		26.83***
		(14.14)		(9.846)
Constant	678.0***	643.8***	714.0***	712.4***
	(53.38)	(5.200)	(0.272)	(0.111)
R-Squared	0.415	0.414	0.604	0.604
Outcome Mean	680.5	680.5	712.78	712.78
Sample	Induced Switches	Induced Switches	All Switches	All Switches
Observations	41279	41279	1343678	1343678

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.1. Effects of Quality and Cost on Spending

Notes: This table shows the impacts of switching to a single point higher quality or lower cost doctor on total spending. Column 1 displays the coefficients on relative switch year indicators, while column 2 displays the impacts from a specification which collapses relative year indicators into a single post-switch indicator. Columns 3 and 4 repeat these analyses on the full sample which includes endogenous switches.

spending because the cost score increases while the quality score (which leads to higher spending) increases. This is also true as shown in Appendix Figure A.3, where the bars for a switch that results in higher quality and cost scores are the highest for both OLS and CS. On the other hand, a switch that results in lower quality and cost scores should lead to the smallest increase in spending. Indeed, as seen in Appendix Figure A.3, bars for these switches (shown at the far right of the figure) are the lowest.

While the literature is still finalizing estimation techniques to allow for robust estimation of the two-way fixed-effects design in this study, these checks lead me to conclude that treatment time heterogeneity may attenuate my OLS estimates while preserving the relative impacts of patient switches across physicians of various quality and cost.

CPT Code	Description
99201	New patient outpatient visit, low complexity, low severity
99202	New patient outpatient visit, low complexity, low to moderate severity
99203	New patient outpatient visit, low complexity, moderate severity
99204	New patient outpatient visit, moderate complexity, moderate to high severity
99205	New patient outpatient visit, high complexity, moderate to high severity
99381	Initial comprehensive preventative medicine evaluation for a new patient: infant
99382	Initial comprehensive preventative medicine evaluation for a new patient: 1-4 years
99383	Initial comprehensive preventative medicine evaluation for a new patient: 5-11 years
99384	Initial comprehensive preventative medicine evaluation for a new patient: 12-17 years
99385	Initial comprehensive preventative medicine evaluation for a new patient: 18-39 years
99386	Initial comprehensive preventative medicine evaluation for a new patient: 40-46 years
99387	Initial comprehensive preventative medicine evaluation for a new patient: 65 years and older
92004	Opthalomogical services: new patient
92002	Opthalomogical services: new patient with diagnostic treatment program

Table A.2. Procedure Codes for New Patients

Notes: This table lists the set of procedure codes used to identify new patient visits.

	(1)
	Number new patients
Female	-0.215
	(-0.54)
Constant	3.712***
	(20.15)
Observations	1587
R-Squared	0.489

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3. Predicting New Patients Using Provider Characteristics

Notes: This table shows the results of a regression which predicts new patients based on physician gender with ZIP code fixed effects. The table uses a cross-section of primary care provider visits from the three months before providers were notified of their new scores.

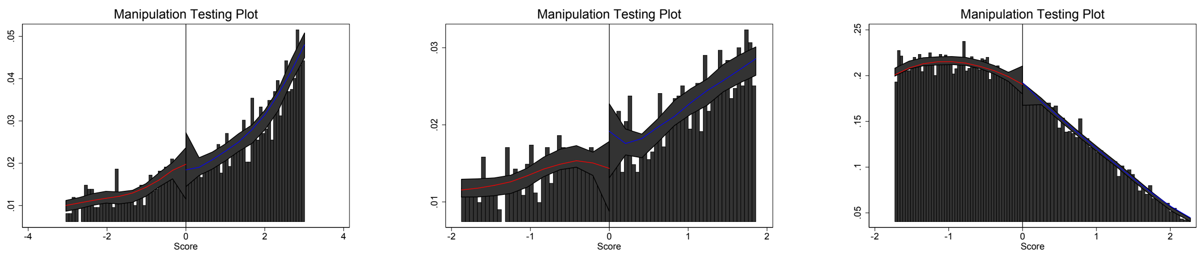

(a) Lower Quality vs. Higher Quality, Lower Cost

(b) Lower Quality vs. Higher Quality, Not Lower Cost

(c) Higher Quality, Lower Cost vs. Higher Quality, Not Lower Cost

Figure A.4. Density Tests

Notes: These figures display the output of density testing from Calonico et al. (2014)'s procedure. There are no statistically significant discontinuities in density across the three designs.

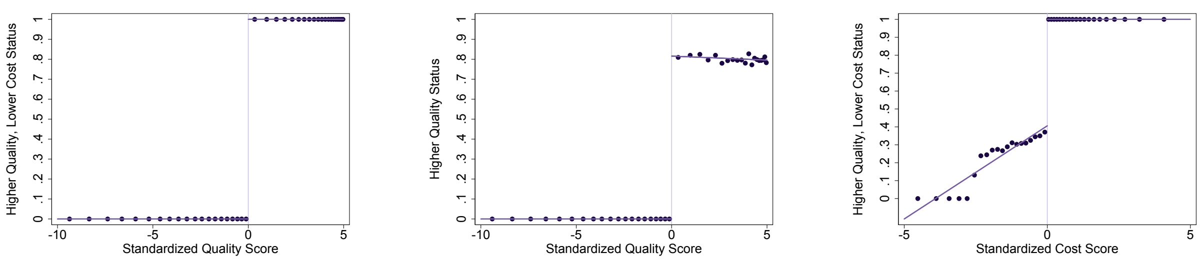

(a) Lower Quality vs. Higher Quality, Lower Cost

(b) Lower Quality vs. Higher Quality, Not Lower Cost

(c) Higher Quality, Lower Cost vs. Higher Quality, Not Lower Cost

Figure A.5. First Stage

Notes: The proportion of "higher quality" (not lower cost) or "higher quality, lower cost" status physicians in each bin for twenty quantile-spaced bins around the cutoff is plotted over bins in Panels A-C. Panel A shows the first stage impact on status for lower-cost physicians, panel B shows the first stage for not lower cost physicians, Panel C shows the impact for higher-quality physicians. Compliance is not perfect in panels B and C since some physicians who would have had quality status are bumped up to "higher quality, lower cost" status if their practice group is considered lower-cost, even if they individually are not. This is only the case for cost criteria: physicians who do not meet quality criteria are always classified as "lower-quality" regardless of their group's behavior.

	(1)	(2)	(3)
	Higher Quality, Lower Cost	Number new patients	Number new patients
Panel A: Quality Fuzzy RD			
Above Quality Cutoff	0.815***	0.230	
	(0.0113)	(0.286)	
Higher Quality Status			0.282
			(0.351)
Regression Type	First Stage	Reduced Form	IV Regression
F	6446.9	3.722	
Observations	19403	19403	19403
R^2	0.499	0.000553	0.000469
	(1)	(2)	(3)
Panel B: Cost Fuzzy RD			
	Higher Quality, Lower Cost	Number new patients	Number new patients
Below Cost Cutoff	0.660***	0.399	
	(0.0141)	(0.306)	
Higher Quality, Lower Cost Status			0.605
			(0.464)
Regression Type	First Stage	Reduced Form	IV Regression
F	.	3.486	
N	7778	7778	7778
R^2	0.514	0.00139	0.000280

Table A.4. IV Estimates

Notes: This table displays instrumental variables estimates from a fuzzy RD design where crossing the threshold is an instrument for having "higher quality" or "higher quality, lower cost" status.

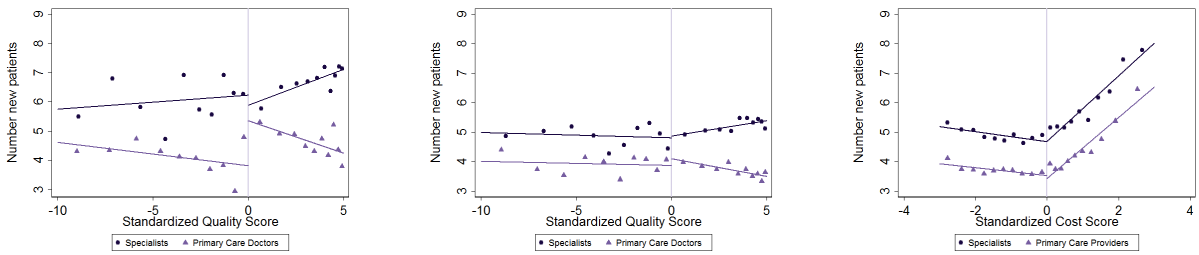

(a) Lower Quality vs. Higher Quality, Lower Cost

(b) Lower Quality vs. Higher Quality, Not Lower Cost

(c) Higher Quality, Lower Cost vs. Higher Quality, Not Lower Cost

Figure A.6. These figures plot the number of new patients seen by providers over the three months following status updates separately for primary care providers and specialists. Outcomes are averaged in each bin for twenty quantile-spaced bins around the cutoff. Panel A shows the impact on new patients for lower-cost primary care physicians. These providers achieved "lower quality" status if below the cutoff, and "higher quality, lower cost" if above. Panel B shows the impact on new patients for non-lower-cost physicians (comparing "higher quality" (not lower cost) status to "lower-quality"), and Panel C shows the impact on new patients for higher-quality physicians based on whether the physician cost score was below the cost threshold (comparing "higher quality, lower cost" to "higher quality" (not lower cost)).

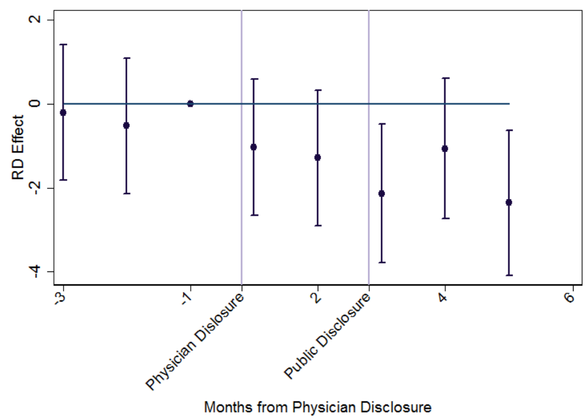

(a) Age of New Patients

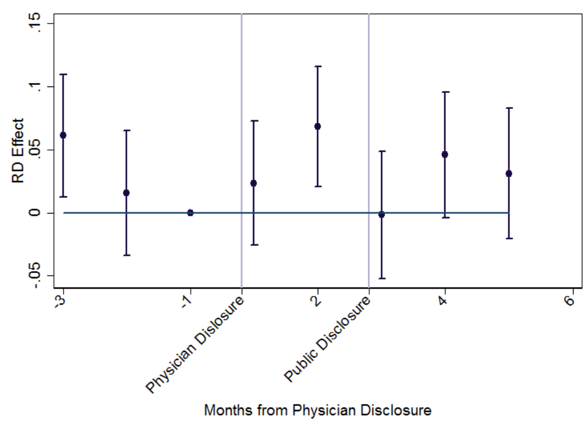

(b) Prop. Chronically Ill of New Patients

Figure A.7. Difference in Discontinuity on Patient Characteristics
Notes: These figures plot the average age and proportion chronically ill new patients of providers over the three months following status updates for lower-cost primary care providers. Panel A displays average age of new patients, and Panel B displays proportion chronically ill. Outcomes are averaged in each bin for twenty quantile-spaced bins around the cutoff.

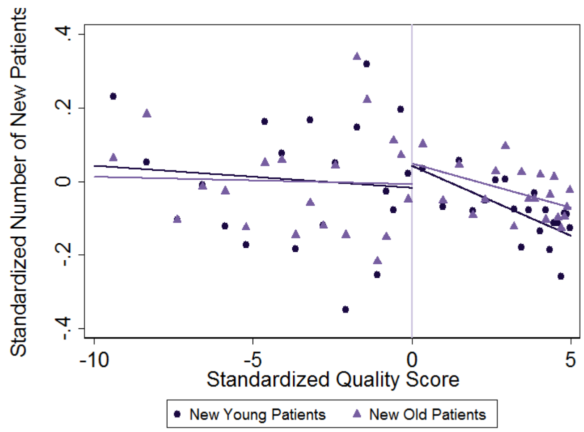

(a) Younger Versus Older Patients

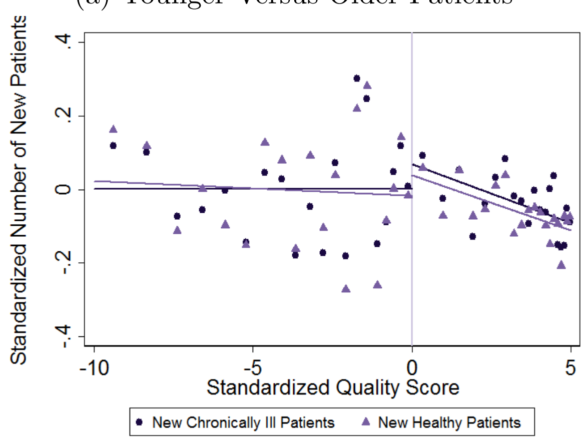

(b) Chronically Ill Versus Healthy Patients

Figure A.8. These figures plot the number of new patients seen by not lower-cost providers over the three months following status updates for primary care providers. New patients are broken down into older (age 40 and above) and younger (age 18-39) patients in panel A and into chronically ill versus healthy patients in panel B. Outcomes are averaged in each bin for twenty quantile-spaced bins around the cutoff.

Dataset	Number of unique patients
Patients who have at least one medical claim each year	4,699,468
Patients whose modal primary care physician is eligible for a quality/cost score	2,657,447
Patients who have only one modal primary care physician	2,618,109
Patients who see a primary care physician each year	197,923
Patients with induced switches	41,972

Table A.5. Switcher Analysis Dataset Sampling

Notes: This table displays the number of unique patients in various subsets of claims data used in the switcher analysis. Subsetting to patients who interact with the medical system once per year removes a large portion of patients. Other subsetting decisions are more marginal.

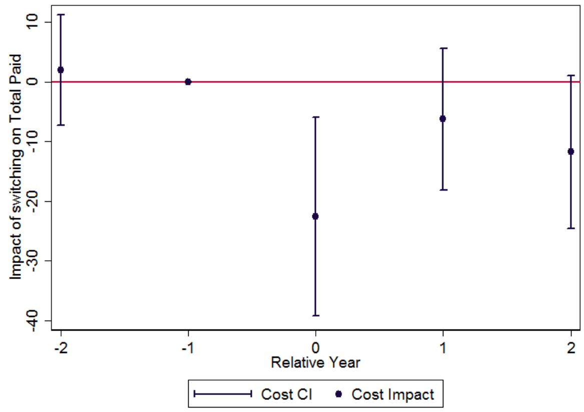

(a) Effect of Cost

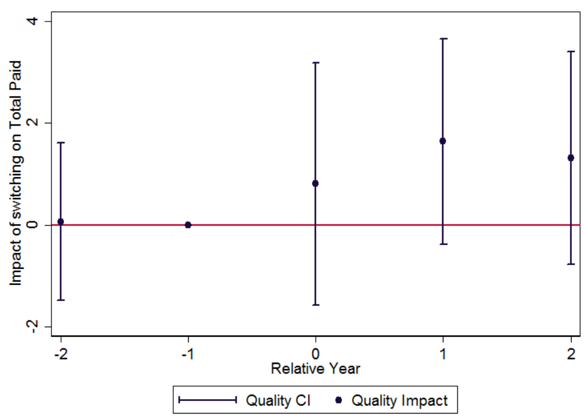

(b) Effect of Quality

Figure A.9. Effect of quality and cost scores on total amount paid
Notes: This figure displays the impacts of switching to a single point lower cost (panel A) or higher quality (panel B) physician on the total amount paid: the sum of patient- and insurer spending estimated on a panel which is balanced on years relative to the switch year.

	(1) Total Paid	(2) Services	(3) Services per Visit	(4) Price per service	(5) Total Paid All	(6) Non-modal Doctor Paid	(7) Patient Paid
$\Delta Q\times$ Post	1.587**	0.0175**	0.00506***	-0.169**	16.20	14.61	0.391
	(0.640)	(0.00719)	(0.00174)	(0.0814)	(12.69)	(12.68)	(0.302)
$\Delta C\times$ Post	-11.49***	-0.0113	0.00932	-1.800***	-107.9**	-96.37**	-2.822*
	(3.660)	(0.0420)	(0.00826)	(0.517)	(47.35)	(47.10)	(1.550)
$\Delta(Q \times C)\times$ Post	0.0889	0.000397	0.000459	0.0731**	0.248	0.159	-0.0784
	(0.353)	(0.00253)	(0.000744)	(0.0316)	(3.875)	(3.910)	(0.204)
R-Squared	0.415	0.510	0.508	0.547	0.452	0.450	0.453
Outcome Mean	680.5	6.34	2.2	142.58	5553.24	4872.74	4872.74
Average Impact	-16.68	.02	.04	-3	-158.88	-142.22	-4.96
Observations	41279	41279	41279	41279	41279	41279	41279

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.6. Effects of Quality and Cost on Spending Outcomes

Notes: This table shows the impacts of switching to a single point higher quality or lower cost doctor on spending outcomes. Column 1 reports the impacts on total spending: the amount paid for visits with the modal primary care physician by both the patient and the insurer. Column 2 reports impacts on the number of service received by the modal primary care physician. The outcome in column 3 is the average number of services received per visit, and column 4 reports the average price per service: total amount paid divided by total number of services. Columns 5 and 6 explore possible spillovers in spending on other providers. The outcome for the regression in column 5 is spending over all providers, not just the modal primary care provider. Column 6 narrows down to spending over other providers, not including the modal primary care provider. Column 7 displays impacts on the portion the patient paid out of pocket. The specification includes an interaction term between the negative cost scores and quality scores to allow for complementarities.

	(1) Total Paid	(2) Total Paid	(3) Total Paid	(4) Total Paid	(5) Total Paid	(6) Total Paid
$\Delta Q\times$ Post	0.706	2.071**	2.775	0.157	0.460	-0.788
	(0.876)	(0.858)	(3.167)	(0.280)	(0.593)	(1.115)
$\Delta C\times$ Post	-13.15**	-8.130*	-30.13***	-25.65***	-15.78***	-15.81***
	(5.469)	(4.584)	(3.912)	(2.682)	(3.806)	(5.678)
Age Range	18-25	26-35	36-45	46-55	56-64	65 and older
Average Impact	-15.08	-7.64	-38.6	-35.82	-21.82	-23.86
R-Squared	0.370	0.418	0.298	0.393	0.264	0.375
Observations	17801	32907	98399	182133	204498	78064

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.7. Effects of Quality and Cost over Age: All Switches

Notes: This table shows the impacts of switching to a single point higher quality or lower cost doctor on spending outcomes for the full sample, which includes endogenous switches.

	(1) Total Paid	(2) Total Paid	(3) Total Paid	(4) Total Paid	(5) Total Paid	(6) Total Paid
$\Delta Q\times$ Post	8.609*	3.793*	4.376	-0.755	1.550	1.039
	(4.712)	(2.262)	(2.914)	(0.950)	(1.502)	(4.157)
$\Delta C\times$ Post	12.78	11.72	-34.20**	-11.00	-12.87	-21.53
	(18.00)	(16.56)	(15.74)	(11.19)	(9.257)	(18.50)
Age Range	18-25	26-35	36-45	46-55	56-64	65 and older
Average Impact	41.46	19.7	-43.02	-16.86	-13.58	-19.9
R-Squared	0.441	0.375	0.467	0.452	0.498	0.361
Observations	460	735	1755	3870	4820	2050

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.8. Effects of Quality and Cost over Age: Relative Time Balanced

Notes: This table shows the impacts of switching to a single point higher quality or lower cost doctor on spending outcomes for the sample which is balanced on time relative to switch. These results are generally under-powered; however, the pattern of results remains.

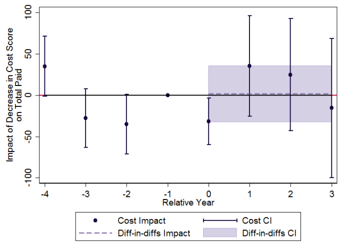

(a) Ages 18-25

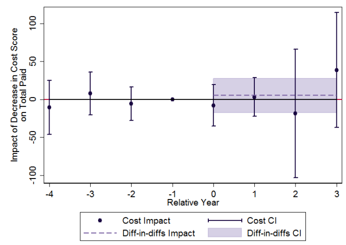

(b) Ages 26-35

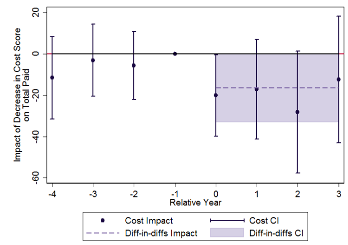

(c) Ages 36-45

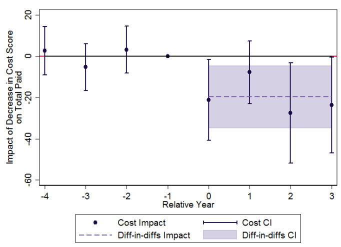

(d) Ages 46-55

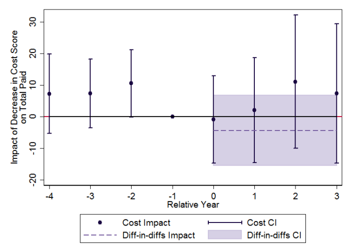

(e) Ages 56-64

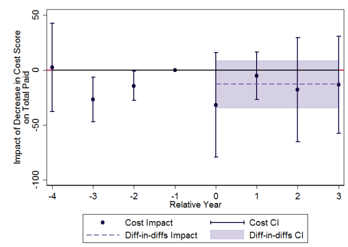

(f) Ages 65 and up

Figure A.10. Effect of quality and cost scores on total amount paid by age group
Notes: This figure displays the impact of cost scores on total amount paid separately by age groups.

	(1) Preventable ED	(2) Preventable ED	(3) Preventable ED	(4) Preventable ED	(5) Preventable ED	(6) Preventable ED
$\Delta Q\times$ Post	0.000146	0.0000386	0.0000259	-0.000134	-0.0000463	0.000106
	(0.000120)	(0.0000918)	(0.0000921)	(0.000103)	(0.000124)	(0.000215)
$\Delta C\times$ Post	0.00180**	-0.00185	0.00116	0.000263	-0.000635	0.000885
	(0.000826)	(0.00127)	(0.000763)	(0.000557)	(0.000764)	(0.00141)
Age Range	18-25	26-35	36-45	46-55	56-64	65 and older
R-Squared	0.515	0.457	0.324	0.359	0.356	0.421
Observations	17801	32907	98399	182133	204498	78064

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.9. Effects of Quality and Cost over Age: All Switches

Notes: This table shows the impacts of switching to a single point higher quality or lower cost doctor on preventable emergency room visits for the full sample, which includes endogenous switches.

	(1) Preventable ED	(2) Preventable ED	(3) Preventable ED	(4) Preventable ED	(5) Preventable ED	(6) Preventable ED
$\Delta Q\times$ Post	-0.000197	0.000854	0.000327	-0.000415*	-0.000229	0.000161
	(0.000649)	(0.00260)	(0.000459)	(0.000232)	(0.000782)	(0.000458)
$\Delta C\times$ Post	0.0114	-0.0215	-0.00391	0.00408	-0.00579**	-0.00763
	(0.00985)	(0.0168)	(0.00254)	(0.00292)	(0.00275)	(0.00581)
Age Range	18-25	26-35	36-45	46-55	56-64	65 and older
R-Squared	0.342	0.497	0.557	0.455	0.387	0.403
Observations	460	735	1755	3870	4820	2050

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.10. Effects of Quality and Cost over Age: Relative Time Balanced

Notes: This table shows the impacts of quality and cost scores on preventable emergency room visits for the sample which is balanced on time relative to switch.

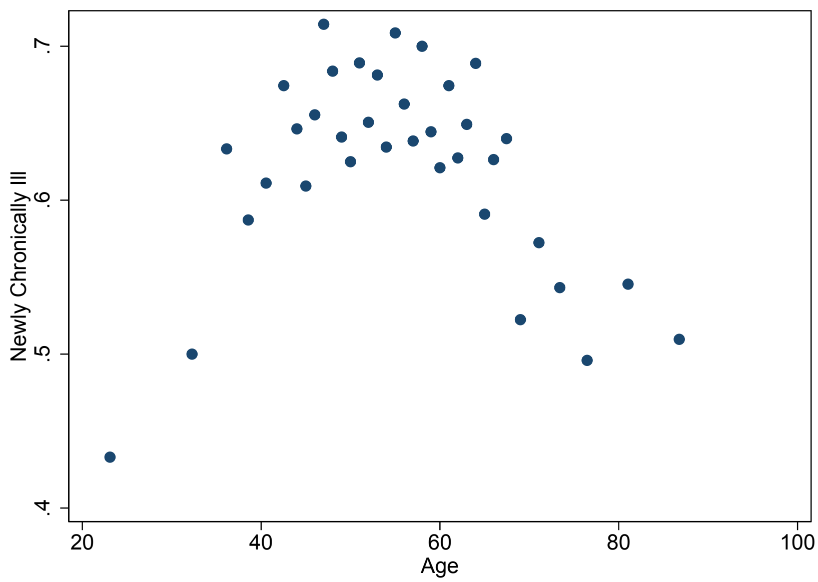

Figure A.11. Binned Scatterplot: Newly Chronically Ill vs Age
Notes: This figure displays a binned scatterplot of the proportion of patients who became chronically ill during the timeframe studies against the age of patients. The patients who were most likely to become chronically ill were middle-aged, between about 40 and 70 years of age.

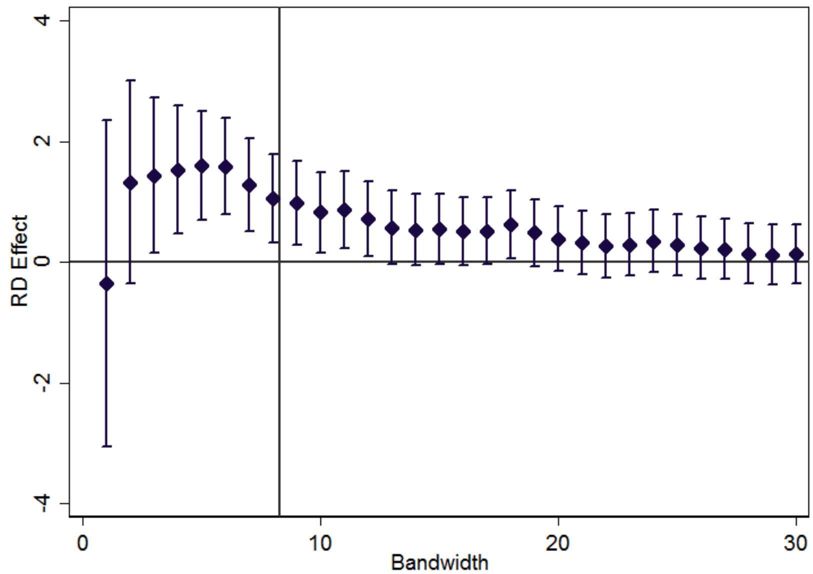

Figure A.12. Results over Multiple Bandwidths
Notes: This figure displays estimates and confidence intervals for the comparison of higher-quality, lower-cost physicians versus lower-quality, lower-cost physicians over multiple bandwidths. For simplicity, bandwidths are the same on either side of the cutoff. Bandwidths from one point to 30 points on either side of the cutoff are shown, and the mean squared error-minimizing bandwidth shown by a vertical line.

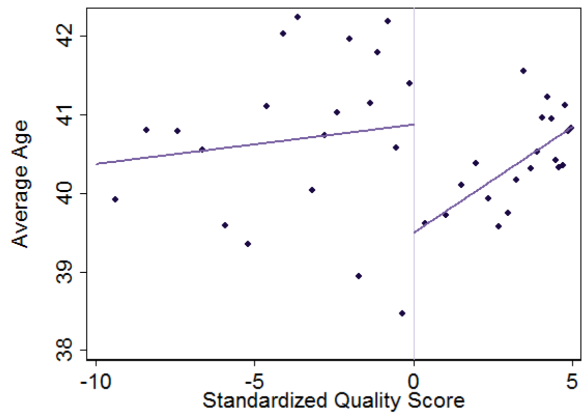

(a) Effect of Passing Quality Cutoff on Patient Age

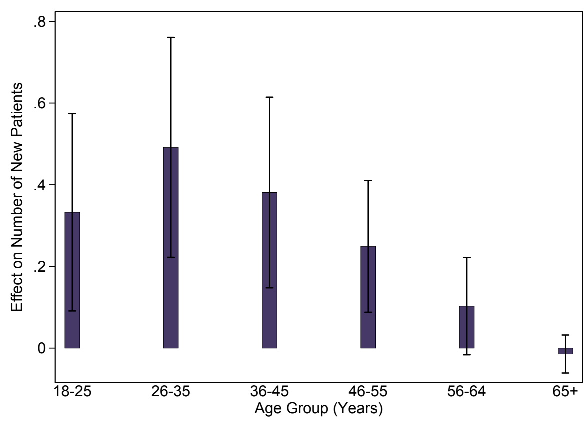

(b) Effect of Passing Quality Cutoff in Age Categories

Figure A.13. Additional Evidence for Age Heterogeneity
Notes: Panel (a) displays the average age of patients treated by lower cost doctors around the Quality threshold, displaying a large negative discontinuity in patient age which suggests that younger patients respond more than older patients. To confirm this, panel (b) shows effects within six age groups, again showing that responses are largest for the younger patients and smaller for the older patients.

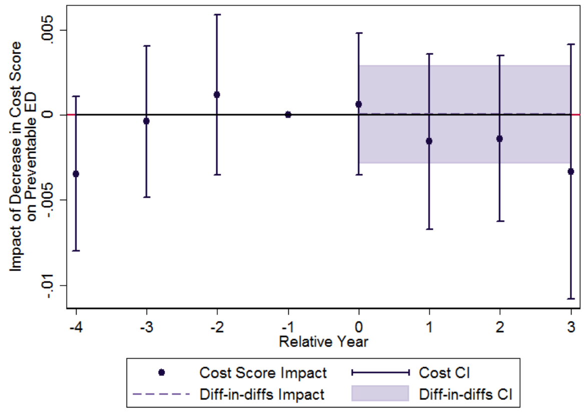

(a) Effect of Cost

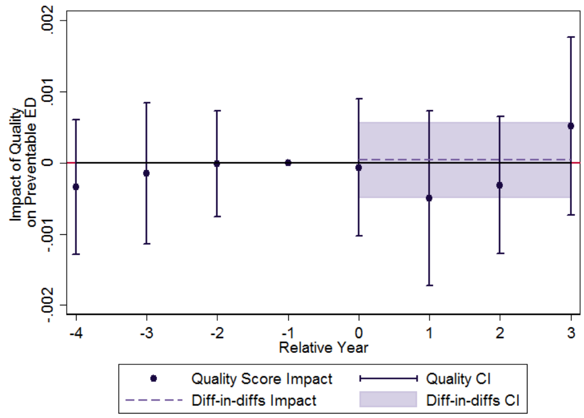

(b) Effect of Quality

Figure A.14. Effect of Cost and Quality on Preventable Emergencies
Notes: Panel (a) displays the impacts of lower cost scores on preventable emergencies, and panel (b) displays the impacts of quality on preventable emergencies.

APPENDIX B

Chapter 2 Appendix

B.1. Treatment Determination Checks

In this section, we present checks to our measure of the number of units in a zip code who have been treated. We want to ensure first that our data cleaning procedures are recovering accurate counts of the overall housing stock and then that we are addressing measurement issues in the number of rent controlled units in a reasonable way.

We first validate our data against the 2000 census in Table B.1. We find only minimal differences, which could be due to buildings being demolished or built between 1990 and 2000.

Table B.1. Comparison of Housing Stock Against US Census

Building Size	Assessor Data - 1999	Census 2000
Single Family Home	118,078	111,125
Two to Four Units	72,646	80,168
Five to Nine Units	34,671	38,940
Ten to Ninteen Units	32,900	34,996
Twenty or More Units	65,838	79,469
Total Units	324,133	344,698

Notes: We construct aggregate measures for all of San Francisco of the number of units that fall in each category of building. Data Sources: 1999 San Fransisco Assessor's Secure Housing Roll and 2000 U.S. Census.

We then construct alternative measures of treatment that attempt to account for various sources of mis-measurement in our primary measure of treatment that largely arise from the fact that the Assessor data is from several years after the policy change. The first concern is that the exemption was only for owner-occupied buildings. We are hesitant to use the owner address data to identify owner-occupied buildings since our

Assessor data is from 1999 several years after the policy, and owner-occupancy may have responded to the policy.[1] For this reason, our preferred measure does not account for this. We construct an alternative measure that does use this information from 1999.

A second concern is that newly rent controlled buildings may have been demolished and replaced with alternative buildings or converted into condos. We construct several measures that would account for this. We alternatively assume that all single family homes build post 1995 replaced rent controlled buildings, that all condos built before 1979 were converted from otherwise rent controlled apartments, that all condos modified in 1995-1999 were condo conversions, that all new condo construction replaced rent controlled buildings, and that all new buildings replaced rent controlled buildings.

Table B.2 reports the correlation between these measures and a falsification measure that is the number of buildings with 5-9 units whose rent control status didn't change. Each of these measures of treatment are very highly correlated, largely because there was not substantial new construction in San Francisco in the 1990s. The least correlated measures with our preferred treatment measure are those that involve older or modified condos. To further convey how similar these various measures are, Figure B.1 shows maps of the various alternative measures of treatment we use.

[1] It is possible that landlords occupied units in their small buildings to be eligible for the rent control exemption prior to 1994; once this exemption was removed, the incentive to live in the building is lower.

Table B.2. Cross-correlation table

Variables	Primary	Owner Occ.	SF	Old Condos	Mod. Condos	New Condos	New Builds	Falsification
Primary	1.000							
Owner Occ.	0.990	1.000						
SF	0.997	0.985	1.000					
Old Condos	0.958	0.968	0.950	1.000				
Mod. Condos	0.949	0.957	0.949	0.988	1.000			
New Condos	0.997	0.987	0.999	0.955	0.954	1.000		
New Builds	0.999	0.989	0.998	0.957	0.950	0.998	1.000	
Falsification	0.878	0.883	0.870	0.929	0.929	0.872	0.877	1.000

Figure B.1. Alternative Treatment Measures

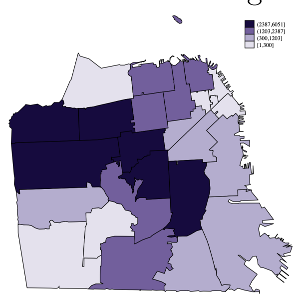

(a) Primary

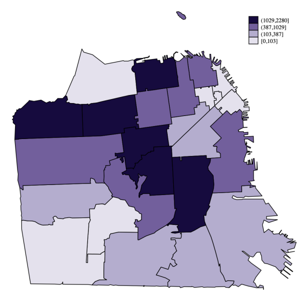

(b) Owner Occupied

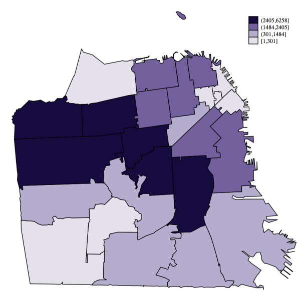

(c) Single Family Adjusted

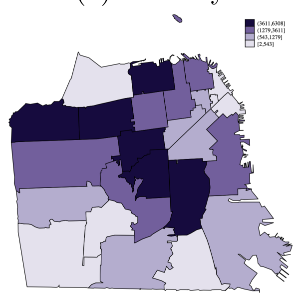

(d) Add Old Condos

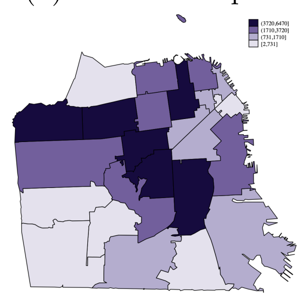

(e) Add Modified Condos

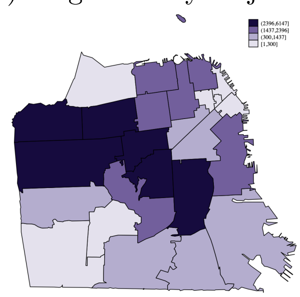

(f) Add New Condos

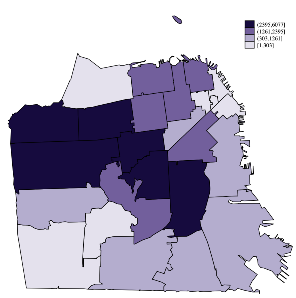

(g) Add New Builds

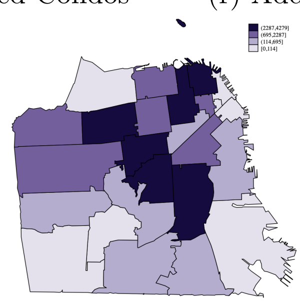

(h) Falsification Test: 5-9 units, built before 1980

Notes: Panel (a) shows a map of our preferred exposure measure that has the number of units in buildings with 2-4 units built before 1980. Panel (b) adjusts this measure for owner occupancy. Panel (c) assumes that all new single family homes built between 1995 and 1999 replaced a duplex. Panel (d) assumes all condos built before 1980 were converted to condos after 1994. Panel (e) assumes all condos modified in the post period were treated. Panel (f) assumes all new condos replaced treated units. Panel (g) assumes all new builds replaced treated units. Panel (h) is our placebo measure of buildings with 5-9 units built before 1980.

While the correlations between these measures are sufficiently high to suggest different treatment measures will not affect our results, we can explicitly test for this. Table B.3 reports coefficient estimates where we include these various measures of treatment as alternatives to our preferred specification. Figure B.3 shows corresponding event study estimates.

Table B.3. Various Measures of Treatment

	(1)	(2)	(3)	(4)	(5)	(6)
Post=1 × IHS NumTreated Owner Occupied	-0.0721* (0.0357)					
Post=1 × IHS NumTreated Single Family Homes		-0.0971*** (0.0308)				
Post=1 × IHS NumTreated Recent Old Condos			-0.0734* (0.0363)			
Post=1 × IHS NumTreated Recent Mod. Condos				-0.0943*** (0.0320)		
Post=1 × IHS NumTreated Recent New Condos					-0.0968*** (0.0308)	
Post=1 × IHS NumTreated New Construction						-0.0843** (0.0332)
Observations	275	275	275	275	275	275
R-Squared	0.811	0.813	0.810	0.811	0.813	0.812

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table assesses the robustness of our main results to alternative ways of measuring the number of treated units in a zip code. Column 1 counts only units which were small, built before 1980, and owner occupied in 1999 based on the address of the owner and address of the unit in the Assessor's Secure Housing Roll. Column 2 includes single family homes as potentially treated units. Column 3 includes condos which were built before 1980, Column 4 includes condos whose last modification date falls within 1994 and 2000, and column 5 includes condos which were built betweeen 1994 and 2000. Column 6 uses a measure of treatment which assumes any building built after 1994 replaced a building that was rent controlled in 1994.

Figure B.3. Event Studies: Alternative Measures of Treatment

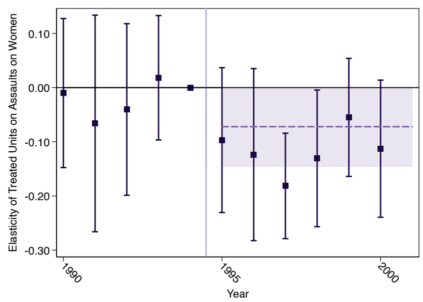

(a) Owner Occupancy Measure

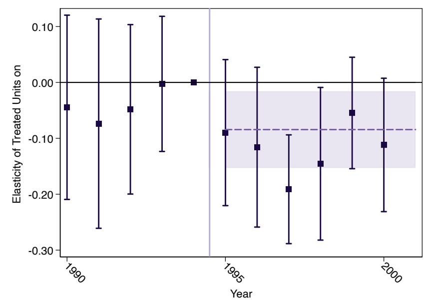

(b) New

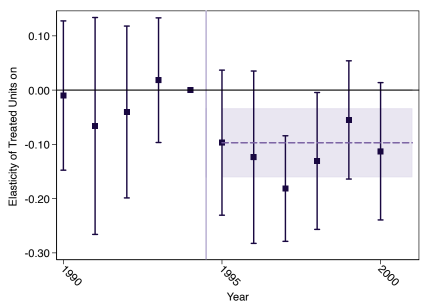

(c) SF

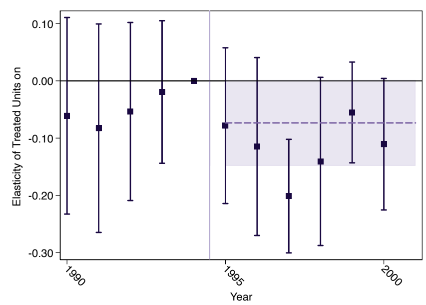

(d) Old Condos

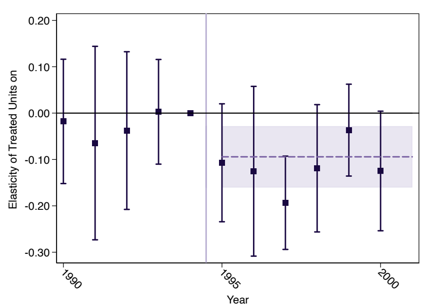

(e) Mod. Condos

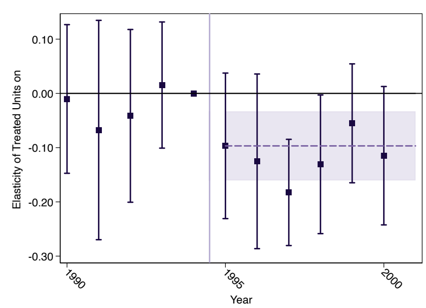

(f) New Condos

Notes: Panels A-F display event study versions of the estimates displayed in Table B.3

We additionally check whether we can predict the number of treated units from Census characteristics in Table B.4. We do not find evidence of strong relationships between observed characteristics of zip codes in 1990 and the number of units that become rent controlled.

Table B.4. Relationship Between Zip Code Census Covariates and Number of Treated Units

	(1) Number of Treated Units
Median Rent	-0.00380
	(0.00403)
Median HH Income	-0.00000598
	(0.0000932)
% Population Black	-2.434
	(3.031)
% Population White	2.775
	(3.294)
% Owner Occupied	1.675
	(2.703)
% Welfare	-5.734
	(9.302)
Constant	4.111
	(3.632)
N	25
R^2	0.129

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: We regress the number of treated units on census demographic variables to determine whether highly treated zip codes were different demographically than less-treated zip codes. Treatment is measured by the number of apartments in the zip code with between two and four units and that was built prior to 1980. Data Sources: 1990 U.S. Census and 1999 San Francisco Assessor's Secure Housing Roll.

B.2. External Cause of Injury Codes

To define intimate partner violence in the hospitalization data, we used external cause of injury diagnostic codes. The codes that we use and their definitions are listed below.

E-codes in the range E960-E969 refer to "Homicide and Injury Purposely Inflicted by Other Persons" and include:

- "Fight, Brawl, Rape"
- "Assault by corrosive or caustic substance, except poisoning"
- "Assault by poisoning"
- "Assault by hanging and strangulation"
- "Assault by submersion [drowning]"
- "Assault by firearms and explosives"
- "Assault by cutting and piercing instrument"
- "Perpetrator of child and adult abuse"
- "Assault by other and unspecified means"
- "Late effects of injury purposely inflicted by other person".

E-codes between 980 and 989 refer to "Injury Undetermined Whether Accidentally or Purposely Inflicted". E904 refers to "Accident due to hunger, thirst, exposure and neglect".

We use the following ICD-9 codes to identify hospitalizations related to substance use disorders or substance abuse:

- 291: "Alcoholic psychoses"
- 292: "Drug psychoses"

- 303: “Alcohol dependency”
- 304: “Drug dependence”
- 305: “Nondependent Abuse of Drugs”

We use the following external cause of injury codes to identify accidents:

- 920: “Unintentional Cut/Pierce”
- 830, 832, 910: “Unintentional Drowning/Submersion”
- 880-889: “Unintentional Fall”
- 890-899: “Unintentional Fire”
- 924: “Unintentional Hot Object”
- 922: “Unintentional Fire Arm”
- 850-869: “Unintentional Poisoning”
- 916-917: “Unintentional Striking”
- 911-913: “Unintentional Suffocation”

B.3. Additional Specifications

In this section, we present results from alternative specifications to show the robustness of our results. Figure C1 shows event study coefficient estimates corresponding to the regression coefficients shown in Table 2.2, in Columns (2)-(5). Our results are largely robust to the inclusion of these additional controls; including the number previously rent controlled makes our estimates noisier but does not substantively change our point estimates in the event study specification.

We additional present event study estimates for alternative ways of specifying our model in Figure C3.

We have now presented robustness to alternative controls, measures of treatment, and model specification. Of course, there are numerous combinations of these checks that we could do. We present a selection of them in Tables C1 and C2.

Figure C1. Event Studies: Alternative Controls

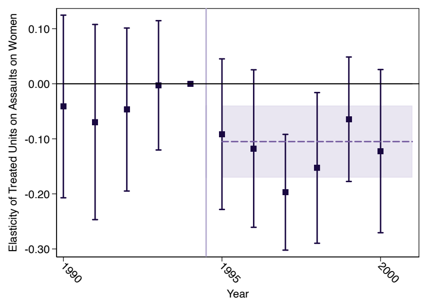

(a) Includes tercile time trends

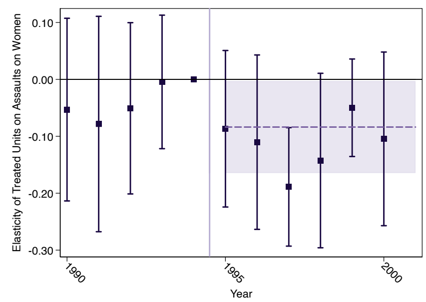

(b) Male assault tercile time trends

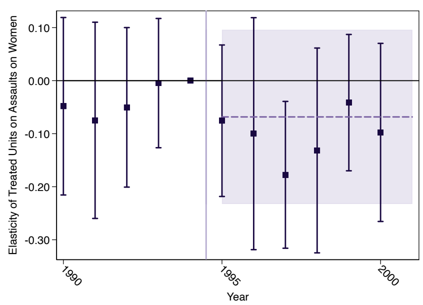

(c) Control for # previously rent controlled

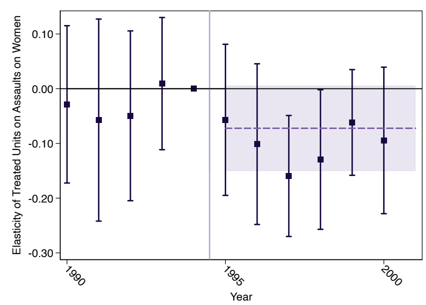

(d) Control for hospital demographics

Notes: Panel (a) shows a specification which controls for treatment tercile-specific linear trends. Panel (b) instead controls for male assault tercile-specific linear trends. Panel (c) includes an interaction term for the number of previously rent controlled units. Panel (d) includes an interaction term with various patient demographic characteristics. Data Sources: OSHPD Inpatient Database and 1999 San Francisco Assessor's Secure Housing Roll.

Figure C3. Event Studies: Alternative Specifications

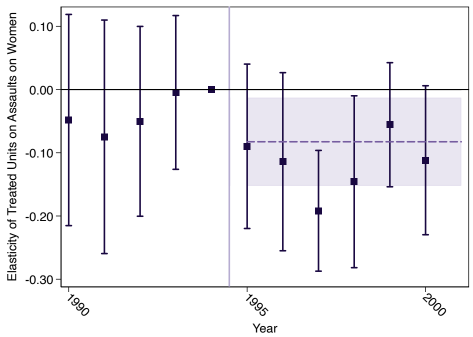

(a) IHS

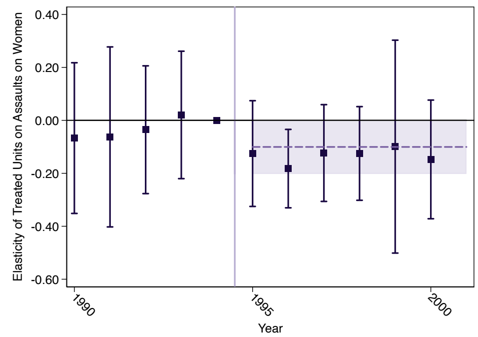

(b) Log

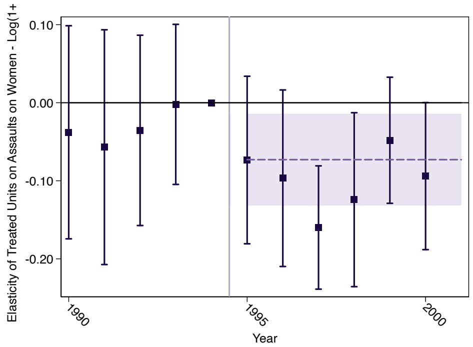

(c) Bottom Coded Log

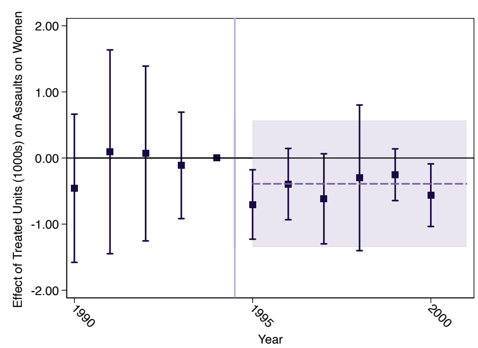

(d) Linear

Notes: Panel (a) shows our specification in logs. Panel (b) shows an alternative specification instead using the inverse hyperbolic sine. Panel (c) replaces all zeros with one to avoid dropping these observations. Panel (d) is a linear specification. Data Sources: OSHPD Inpatient Database and 1999 San Francisco Assessor's Secure Housing Roll.

Table C1. Log DID

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post=1 ×	-0.100*					-0.173***	-0.0819	-0.0999**
Log NumTreated	(0.0493)					(0.0417)	(0.0711)	(0.0466)
Post=1 ×		-0.0874*						
Log Owner Occupied		(0.0453)						
Post=1 ×			-0.117**					
Log Single Family			(0.0561)					
Post=1 ×				-0.103*				
Log New Builds				(0.0498)				
Post=1 ×					-0.0916			
Log Condo					(0.0549)			
Post=1 ×							-0.0441	
Log Prev. RC							(0.0775)	
Specification	Main	Owner Occupied	SF	New	Condo	Tercile Trends	Previous RC	Hospital Covs
Observations	233	231	233	233	233	233	231	233
R-Squared	0.799	0.795	0.799	0.799	0.797	0.815	0.796	0.801

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports robustness of results to various controls in the specification, using logs instead of inverse hyperbolic sine. Data Sources: OSHPD Inpatient Database and 1999 San Francisco Assessor's Secure Housing Roll.

Table C2. LOG 1+X DID

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post=1 × Log(1+NumTreated)	-0.0729**					-0.0949***	-0.0589	-0.0637*
	(0.0285)					(0.0266)	(0.0648)	(0.0309)
Post=1 × Log(1+ Owner Occupied)		-0.0657**						
		(0.0303)						
Post=1 × Log(1 + Single Family)			-0.0854***					
			(0.0264)					
Post=1 × Log(1 + New Build)				-0.0744**				
				(0.0280)				
Post=1 × Log(1 + Condo)					-0.0646**			
					(0.0299)			
Post=1 × Log(1 + Prev RC)							-0.0162	
							(0.0562)	
Specification	Main	Owner Occupied	SF	New	Condo	Tercile Trends	Previous RC	Hospital Covs
Observations	275	275	275	275	275	275	275	275
R-Squared	0.821	0.820	0.822	0.821	0.819	0.826	0.821	0.825

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports robustness of results to various controls in the specification, but adjusting variables using log(1+X) so that variables with with a value of zero can be included in the analysis. instead of inverse hyperbolic sine. Data Sources: OSHPD Inpatient Database and 1999 San Francisco Assessor's Secure Housing Roll.

Finally, we report patterns in the raw data in figure C5, plotting assaults on women for quartiles of treatment exposure over time.

Figure C5. Patterns in Raw Data: Assaults on Women by Quartile of Treatment Exposure

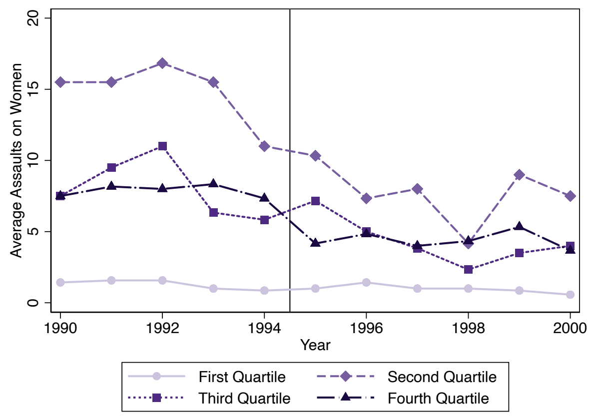

Notes: This figure displays assaults on women over time, plotted by quartile of treatment exposure. The relative decrease in assaults on women is most noticeable in a comparison of the first versus four quartile of exposure. This table reports robustness of results to various controls in the specification, using logs instead of inverse hyperbolic sine. Data Source: OSHPD Inpatient Database and 1999 San Francisco Assessor's Secure Housing Roll.

B.4. Demographic Patterns

One benefit of using hospitalization data is that it provides rich information on the characteristics of hospitalizations, both for assaulted and non-assaulted patients. This information allows us to confirm that our measure of IPV conforms with the expected patterns from the literature and to assess whether there are other changes in hospitalization patterns in response to the policy.

Figure D1. Age Distribution of Hospitalized Women

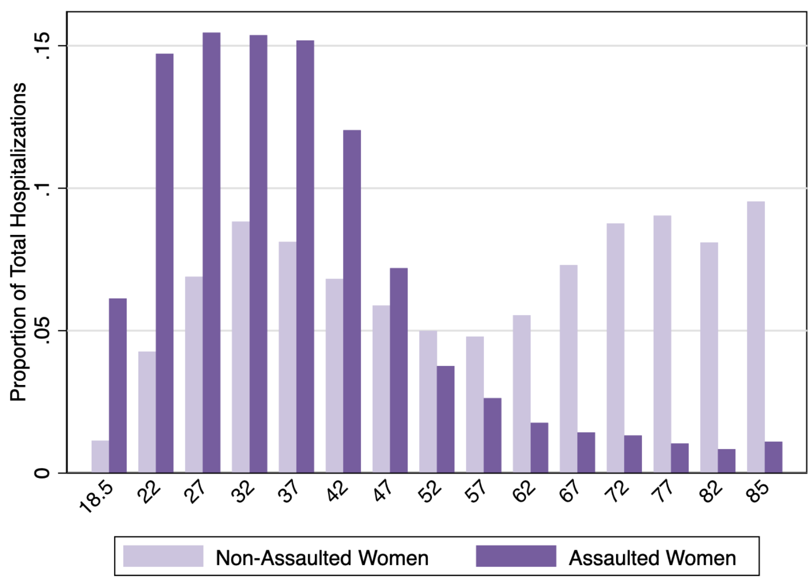

Notes: We display a histogram depicting the age structure of assaulted women. Frequencies for all female patients are displayed in dark gray, while frequencies for assaulted female patients are shown in light gray. Assaulted patients appear more likely to be young. Data Source: OSHPD Hospitalization Data.

Figure D2. Racial Composition of Hospitalized Women

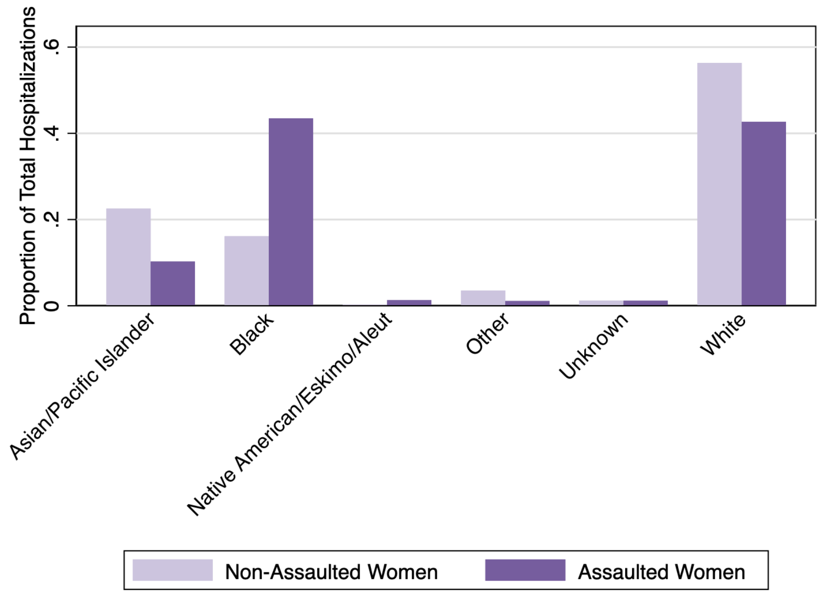

Notes: We display a histogram depicting the race structure of assaulted women. Frequencies for all female patients are shown in dark gray, while frequencies for assaulted female patients are shown in light gray. Assaulted women appear more likely to be black. Data Source: OSHPD Hospitalization Data.

To predict intimate partner violence, we first run a regression of IPV on the number of women who were admitted to the hospital from different race and ethnic categories and the median age of hospitalized women. The results of this predictive equation are shown in Appendix Table D1.

In this section, we present additional event study plots looking at the time evolution of the effects of the variation induced by the policy change on demographic characteristics of who is hospitalized. We presented these results in a difference-in-differences table in the main text.

Table D1. Predicting Assaults on Women Based on Zip Code Demographic Characteristics

	All Covariates
Number Black Patients	0.00639***
	(0.000534)
Number White Patients	0.00127***
	(0.000382)
Number Asian Patients	0.000652
	(0.000630)
Number Hispanic Patients	0.00234***
	(0.000842)
Median Age	-0.118***
	(0.0280)
Constant	6.370***
	(1.427)
R Squared	0.627
Observations	275

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Notes: This table shows the results of a regression using demographic characteristics to predict assaults on women using demographic information from hospitalized patients. Data source: OSHPD

Figure D3. Effects of Rent Control Exposure on Median Age of Female Patients

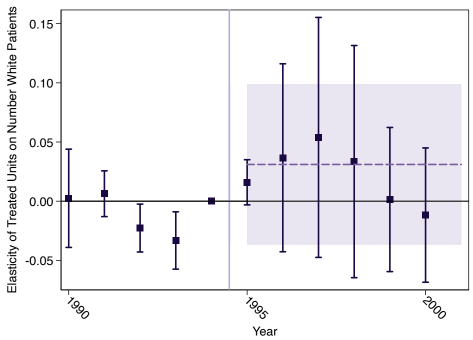

(a) White Patients

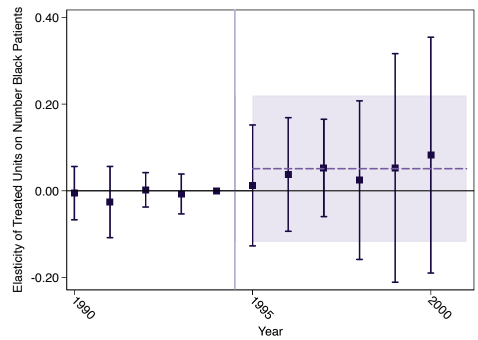

(b) Black Patients

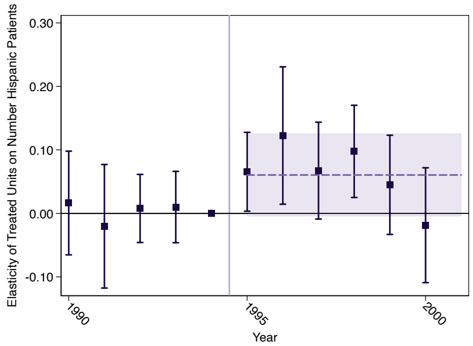

(c) Hispanic Patients

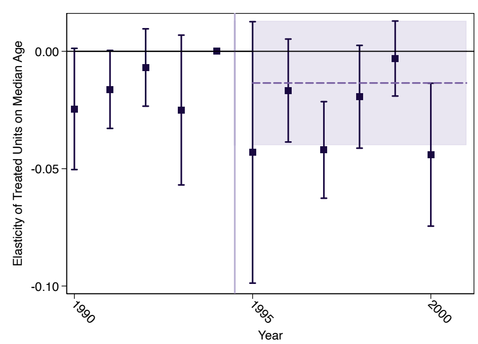

(d) Median Age

Notes: Panel (a) shows the estimated effects for the number of White patients. Panel (b) shows the effects for the number of Black patients. Panel (c) shows the effects for the number of Hispanic patients. Panel (d) shows the effects on the median age of hospitalized women.

B.5. Changes in San Francisco's Housing Market

In this appendix section, we present evidence of rent changes that occur in the San Francisco housing market in response to the policy change. Data on zip code level rent prices dating back to the 1990s is scarce, so we use information on median gross monthly rent from the decennial Census in 1980, 1990, and 2000. These data are self-reported by respondents to the long form Census and so may be variable in whether they include utilities depending on how the respondent interpreted the question.

Figure H1. San Francisco Rent Prices and Rent Control

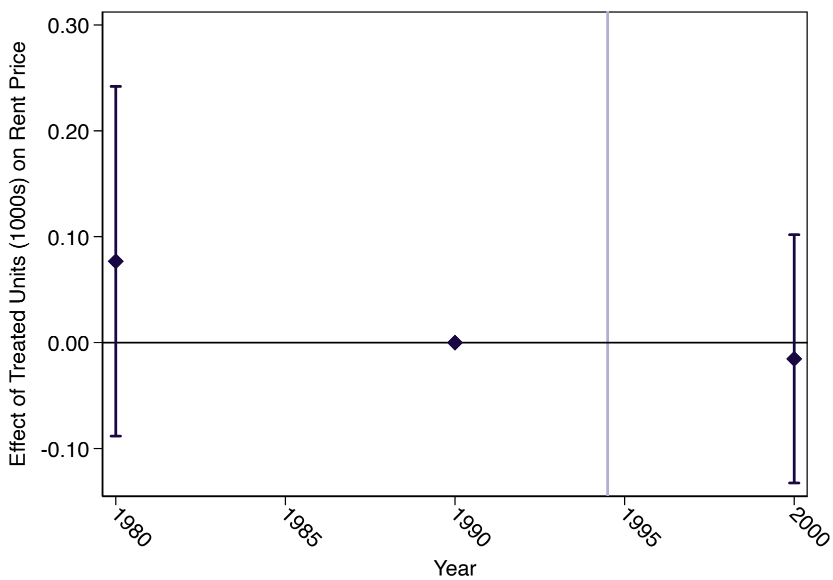

Notes: Event study estimates of the effects of rent control on rent prices. Standard errors are clustered at the zip code level. Error bars show 95% confidence intervals. Data source: 1980, 1990, 2000 US Census.

We estimate an event study specification similar to equation 2.2 where 1990 is our omitted category as the year closest to and previous to the policy change. We find no evidence of pre-trends leading up to 1990 (although we only have two data points so we interpret this finding very cautiously), and then a small decline in rent prices in highly treated zip codes in 2000, though this effect is not statistically significant.

APPENDIX C

Chapter 3 Appendix

C.1. Data Appendix

C.1.1. Assessor Data Cleaning Details

In our analysis, we use data from the San Francisco Assessor's Secure Housing Roll from 1999. The apartment address data from the Assessor's office does not include unit ZIP codes, so we merge the Assessor's data to the San Francisco's Addresses with Units - Enterprise Addressing System dataset (SFA). The SFA includes complete addresses with ZIP codes, using the parcel number or unit address. While we can match most units with this procedure, some are either missing from the SFA data or do not have unit address information in the Assessor's data. For these, we either use the Google Map's API to determine ZIP code based on SITUS or use the San Francisco Planning tool to determine the unit's address based on its parcel number.

To fill in missing fields in the Assessor's dataset, which come from hand-filled forms, we do the following. First, if any building is listed as having zero units, we replace the number of units using the median number of units based on that building's class code. This means that, since Multi-Family Residences with class codes of "Flats and Duplex" have a median of two units, any Multi-Family Residences with that class code will be reassigned to have two units if it has zero units listed in the dataset. Next, for any building that is missing information on the year built, we assume that it was built before 1980 since the median year built of all buildings in San Francisco is 1925.

In Appendix Table A1, we compare the number of rental units in our dataset from 1999 to those in the 2000 census to ensure we have correctly categorized and in some cases, imputed, the building size correctly. The differences between measures from our dataset and the census are quite small, and could be due to buildings being built or demolished between 1999 and 2000.

In Appendix Table A2, we report results of a regression of various ZIP code characteristics on the number of treated units in our sample to show that the number of treated units in each neighborhood was not confounded with neighborhood characteristics. We find no significant differences on any observable characteristics of neighborhoods. While this is not necessary to interpret our difference-in-differences results as causal effects, it helps establish credibility for the required (and untestable) identification assumptions: that neighborhoods with different levels of treatment would have trended similarly if the policy had not passed, and that treatment effects are homogeneous across neighborhoods with different levels of treatment intensity.

C.1.2. Alternative Treatment Measure

While we do not know whether owners occupied their units in 1994 when the policy changed, we can approximate this using a 1999 measure of owner occupancy. To create this, we determine whether the owner's mailing address listed in the Assessor's dataset matches the physical address of the unit. We use this measure of owner occupancy to create an alternative measure of exposure, defined as the number of units in a ZIP code that had between two and four units, were built before 1980, and were owner-occupied in 1999 (the earliest year of Assessor data available). This alternative measure of exposure to the rent control policy is displayed in Appendix Figure A1. While some areas become more or less exposed based on this alternative measure, the distribution of exposure to rent control looks largely similar to that in Appendix Figure 3.1.

In Appendix Figure A1 we display a map of San Francisco, where ZIP codes with higher numbers of treated buildings are shown in darker blue. This figure uses the treatment definition that includes 1999 owner occupancy. In comparing to our preferred measure of treatment which does not account for owner occupancy, it is clear that the two definitions of treatment affect the same neighborhoods to roughly the same extent.

Finally, in Appendix Figure A2, we show the event study results of running our main event study regression using the measure of treatment which includes a 1999 measure of owner occupancy. The pattern of results is the same as those in Figure 3.4, which does not use any measure of owner occupancy to measure treatment status. While the magnitudes of effects are much larger, since the average number of owner-occupied treated units is lower than the measure which does not take owner-occupancy into account, the proportional effects are similar to those in our main specification.

C.1.3. Validating and Cleaning Eviction Notice Data

To measure evictions, we use data from the San Francisco Rent Board, who publishes two measures. First, the Rent Board publishes counts of eviction notices from 1990 and onward for San Francisco as a whole. These counts include eviction notices that were verified by the Rent Board. This dataset is not available at the individual rental unit level before 1997 due to missing fields in the pre-1997 data. However, the total number of annual evictions are published each year starting in 1990. The Rent Board also publishes a ZIPcode-year level dataset of wrongful eviction claims from 1990 and onward. This dataset includes all allegations of wrongful evictions made by tenants, regardless of whether their eviction was legal or not, and regardless of whether the eviction went through.

The Rent Board has provided us with their pre-1997 individual-level eviction notices data with the caveat that the data are missing fields and have not been audited or checked. First, we check that the pre-1997 data are not missing observations. To do this, we compare the annual number of evictions in the pre-1997 + post-1997 dataset to the published annual number of evictions, finding that the numbers match between the pre-1997 data and the published statistics. If substantial numbers of observations were missing in the data, we would expect to see fewer evictions in the pre-1997 data. Since we do not see this, we proceed assuming all evictions are accounted for in the data.

Next, we evaluate missing fields in the pre-1997 dataset. We find that 32% of fields are missing ZIP codes, and 30% of fields are missing the reason for the eviction. Only 0.06% of fields are missing an address. For those observations that are missing ZIP codes, we use Google Maps Places API, searching for the unit's address, and filling in the ZIP code that Google Maps associates with the address. We find ZIP codes for all but 124 of the 2,649 rows with missing ZIP codes, so that overall we have ZIP codes for 98% of the data. To ensure the Google Maps API accurately assigns ZIP codes, we compare Google Maps ZIP codes to the ZIP codes included in the data. We find that 88% of ZIP codes are correctly identified by the Google Maps API. We proceed by using the ZIP code given by the eviction notices data unless it is missing, in which case we use ZIP code given by the Google Maps API, acknowledging the existence of some measurement error in the ZIP code variable. Since there is no way to fill in missing cause of eviction, we proceed with the understanding that eviction cause data may be unreliable.

C.2. Additional Robustness Checks

We report alternative specifications that evaluate the sensitivity of our results to how we measure treatment in Appendix Tables A3 and A4. Columns 1 and 2 are the same as the main specifications reported in the paper. Columns 3 and 4 use a measure of treatment that accounts for owner occupancy. Columns 5 and 6 instead use a binary measure of treatment based on whether ZIP codes are above or below the median level of treatment.

Since only those in rent controlled units were eligible to submit eviction notices and wrongful eviction claims to the Rent Board, we worry that the effect of rent control on evictions is not due to a behavioral change, but instead mechanical. Suppose 1% of tenants are evicted regardless of rent control status. If an additional 100 buildings are subjected to rent control, then an increase of 1 eviction would be expected, mechanically. To explore whether the effects we find are mechanical or due to behavioral change, we examine the proportion of rent controlled tenants that face eviction notices before and after the policy change. If the rate of eviction stays constant over time, then the effects are purely mechanical, whereas if the rate increases at the time of the policy change, then effects are likely due to behavioral changes by landlords. In Appendix Figure A6, we display the number of eviction notices per number of rent controlled units, which increases after 1994, when the policy was passed.

Appendix Figure A6 suggests that just after rent control expands, landlords adjust by evicting tenants, many times wrongfully. The gradual increase after the policy change makes sense; as time passes from the policy change, the difference between market rent, which can freely increase, and rent controlled rents, which can increase by at most about

2% each year, increases. With it increases the incentive to evict a tenant in order to reset the rent to market rate or convert the building to condos. Indeed, Diamond et al. (2019) find an 8 percentage point increase in condo conversions following this same rent control expansion.

C.3. Appendix Figures

Figure A1. ZIP Code Level Exposure to Change in Rent Control Policy: Owner Occupancy

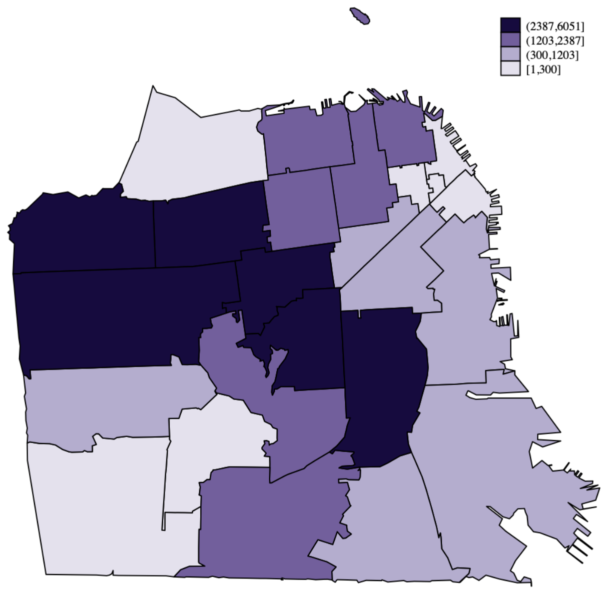

Notes: This map depicts exposure based on owner-occupancy at the ZIP code level for the San Francisco policy change. Exposure is measured as the number of units in a ZIP code in buildings with 2-4 units that were built prior to 1979 where the owners' address matches that of the building.
Data source: 1999 San Francisco Assessor's Secure Housing Roll and authors' calculations.

Figure A2. Effect of Removing Rent Control Exemptions on Eviction Notices and Wrongful Eviction Claims: Owner Occupancy Robustness

(a) Effect on Eviction Notices

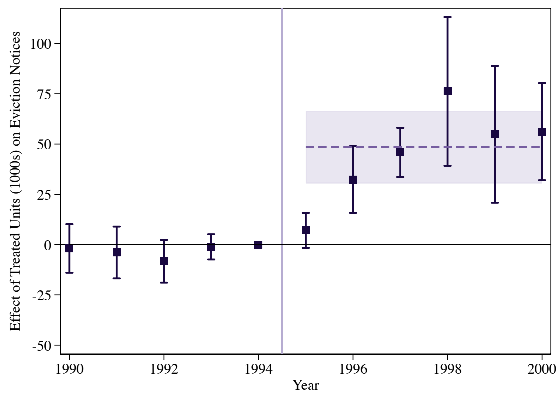

(b) Effect on Wrongful Evictions

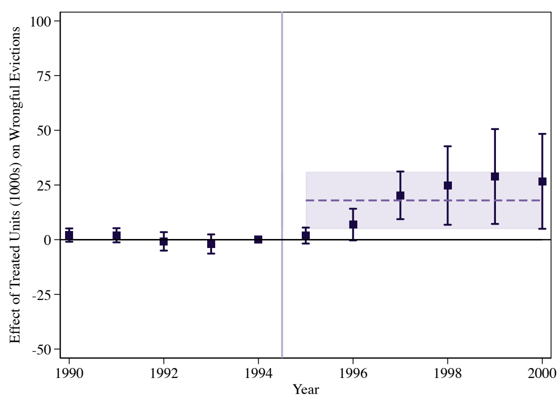

Notes: This figure shows the robustness of event study coefficient estimates to an alternative measure of treatment where we attempt to condition treatment on the building being owner occupied. We show our estimated coefficients on the interaction between year dummies and the number of treated units in a ZIP code. We normalize 1994 to be zero. Error bars shown are for the 95% confidence interval. The difference in difference estimate for the interaction of the post period with the number of treated units is shown as a grey dashed line with the 95% confidence interval shown as a shaded region on the graph. Standard errors are clustered at the ZIP code level. Panel A reports effects on eviction notices, while Panel B reports effects on wrongful eviction claims. Effects are scaled to be the effect per 1,000 treated units. The average number of treated units in a ZIP code is 1,688.

Figure A4. Eviction Notices Over Time

(a) Published Data

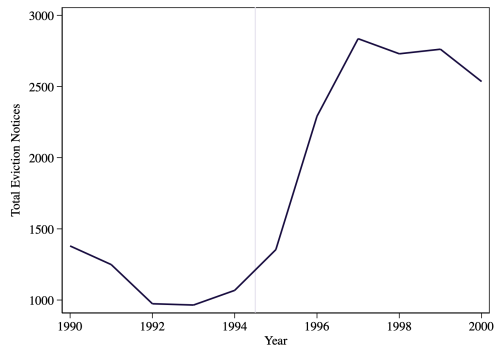

(b) Unpublished Individual-Level Data

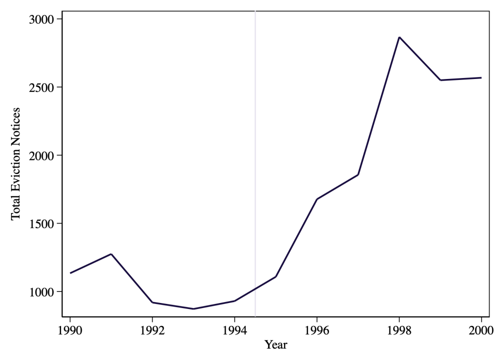

Notes: These figures shows the total number of evictions notices by year. Panel A displays the overall number of eviction notices over time published by the Rent Board, while Panel B reports the same measure based on the unpublished individual-level data.
Data Sources: San Francisco Rent Board Eviction Notices, published and un-published datasets

Figure A6. Number of Evictions per Number of Rent Controlled Units

(a) Eviction Notices

(b) Wrongful Eviction Claims

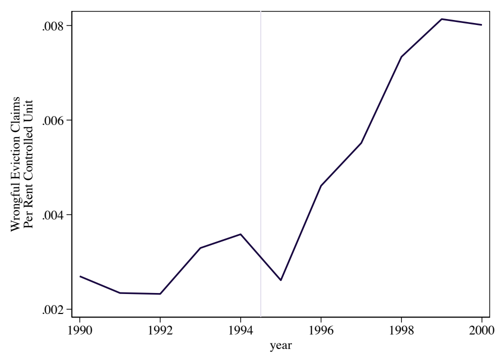

Notes: These figures plot the average fraction of eviction notices (Panel A) and wrongful eviction claims (Panel B) as a fraction of the number of rent controlled units in a given ZIP code over time.
Data Sources: 1999 San Francisco Assessor's Secure Housing Roll and San Francisco Rent Board Eviction Notices and Wrongful Eviction Claims

Figure A8. Effect on Rent Prices

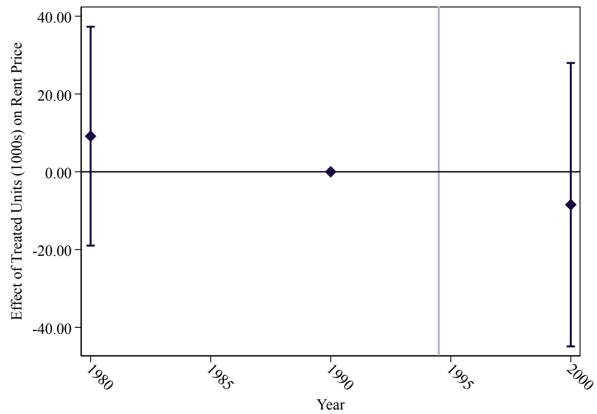

Notes: This figure shows our event study coefficient estimates on the interaction between year dummies and the number of treated units in a ZIP code on self-reported rents from the 1980, 1990, and 2000 decennial census. We normalize 1990 to be zero. Error bars shown are for the 95% confidence interval. Effects are scaled to be the effect per 1,000 treated units. The average number of treated units in a ZIP code is 1,688.
Data Sources: 1999 San Francisco Assessor's Secure Housing Roll and 1980-2000 U.S. Census

Figure A9. Wrongful Eviction Claims Per Eviction Notice Over Time by Tercile of Policy Exposure

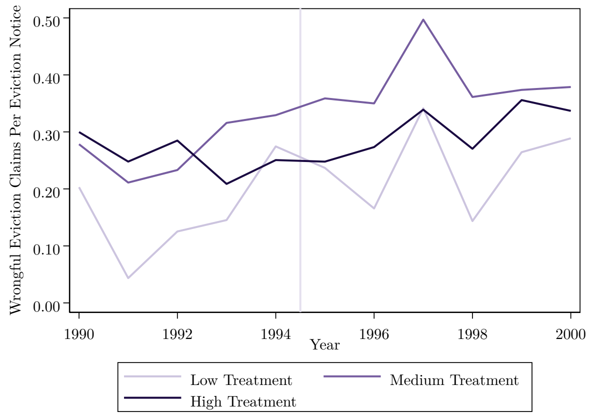

Notes: This figure shows the average fraction of wrongful evictions per eviction notice in ZIPcodes, broken into terciles of policy exposure. We assign a value of zero for ZIP codes with no eviction notices. Low treatment terciles are those in the lowest tercile of units newly exposed to rent control after the policy change. Middle treatement ZIP codes are those in the middle tercile. High treatment are those in the highest tercile. The timing of the policy change is marked by a vertical line.
Data Sources: 1999 San Francisco Assessor's Secure Housing Roll, San Francisco Rent Board Eviction Notices and Wrongful Evictions Claims

Figure A10. Long Run Eviction Notices By Type

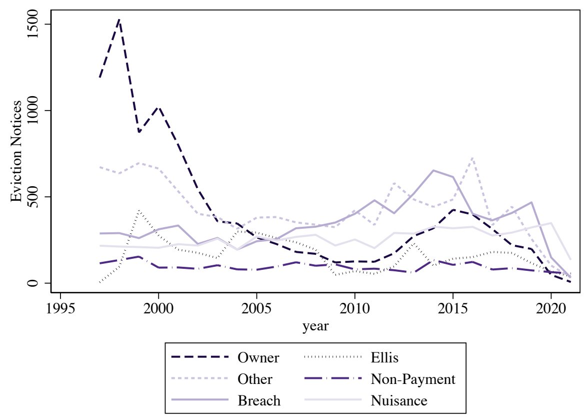

Notes: This figure plots the time series of different types of eviction notices over the long run.
Data Sources: San Francisco Rent Board Eviction Notices

C.4. Appendix Tables

Table A1. Comparison of Housing Stock Against US Census

Building Size	Assessor Data - 1999	Census 2000
Single Family Home	118,078	111,125
Two to Four Units	72,646	80,168
Five to Nine Units	34,671	38,940
Ten to Ninteen Units	32,900	34,996
Twenty or More Units	65,838	79,469
Total Units	324,133	344,698

Notes: We construct aggregate measures for all of San Francisco of the number of units that fall in each category of building and compare them to the same measures from the 2000 Decennial Census.
Data Sources: 1999 San Francisco Assessor's Secure Housing Roll and 2000 U.S. Census.

Table A2. Relationship Between ZIP Code Census Covariates and Number of Treated Units

	(1) Number of Treated Units
Median Rent	-0.00380
	(0.00403)
Median HH Income	-0.00000598
	(0.0000932)
% Population Black	-2.434
	(3.031)
% Population White	2.775
	(3.294)
% Owner Occupied	1.675
	(2.703)
% Welfare	-5.734
	(9.302)
Constant	4.111
	(3.632)
N	25
R^2	0.129

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: We regress the number of treated units on census demographic variables to determine whether highly treated ZIP codes were different demographically than less-treated ZIP codes. Treatment is measured by the number of apartments in the ZIP code with between two and four units and that was built prior to 1980. Specifically, we regress Num $\text{Treated}_j = \beta_0 + \beta_1 1990\ \text{rent}_j + \beta_2 \text{Income}_j + \beta_3 \text{Prop Black}_j + \beta_4 \text{Prop White}_j + \beta_5 \text{Prop Owners}_j + \beta_6 \text{Prop on Welfare}_j + \epsilon_j$ for characteristics of ZIP code j.
Data Sources: 1990 U.S. Census and 1999 San Francisco Assessor's Secure Housing Roll.

Table A3. Difference-in-Difference Estimates of the Effects of Rent Control on Eviction Notices

	(1)	(2)	(3)	(4)	(5)	(6)
Num Treated × Post	20.07***	15.24***				
	(2.926)	(2.374)				
Num Treated × Post × Low Income		9.662***				
		(2.651)				
Owner Occupied × Post			48.45***	34.24***		
			(8.741)	(6.691)		
Owner Occupied × Post × Low Income				30.09***		
				(7.780)		
Treated × Post					51.84***	30.46***
					(11.55)	(10.38)
Treated × Post × Low Income						47.97***
						(9.499)
N	275	275	275	275	275	275
R^2	0.882	0.892	0.877	0.889	0.853	0.883
Wild Bootstrap P-Value	0	0.00700	0	0.0100	0	0.00600
Wild Bootstrap P-Value	.	0.0160	.	0.0170	.	0.0860

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates from difference-in-difference regressions on eviction notices. Column 1 shows the results from our preferred specification, a difference-in-differences regression using the number of treated units at the ZIP code level as a continuous treatment measure. Column 2 adds heterogeneity by median income. Columns 3 and 4 replace the treatment measure with the number of units eligible for the rent control policy that were also owner occupied in 1999. Columns 5 and 6 replace the treatment variable with an indicator for the number of treated units in the ZIP code exceeding the median number of treated units across ZIP codes.

Table A4. Difference-in-Difference Estimates of the Effects of Rent Control on Wrongful Eviction Claims

	(1)	(2)	(3)	(4)	(5)	(6)
Num Treated × Post	7.632***	4.191***				
	(2.344)	(1.184)				
Num Treated × Post × Low Income		6.718**				
		(2.705)				
Owner Occupied × Post			17.95***	8.736**		
			(6.343)	(3.376)		
Owner Occupied × Post × Low Income				19.12**		
				(7.642)		
Treated × Post					20.19***	10.04*
					(5.941)	(5.030)
Treated × Post × Low Income						22.46**
						(8.851)
N	275	275	275	275	275	275
R^2	0.830	0.847	0.822	0.843	0.805	0.843
Wild Bootstrap P-Value	0	0.0140	0.00200	0.0310	0	0.0640
Wild Bootstrap P-Value	.	0.166	.	0.123	.	0.172

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates from difference-in-difference regressions on wrongful eviction claims. Column 1 shows the results from our preferred specification, a difference-in-differences regression using the number of treated units at the ZIP code level as a continuous treatment measure. Column 2 adds heterogeneity by median income. Columns 3 and 4 replace the treatment measure with the number of units eligible for the rent control policy that were also owner occupied in 1999. Columns 5 and 6 replace the treatment variable with an indicator for the number of treated units in the ZIP code exceeding the median number of treated units across ZIP codes.

Table A5. Robustness to Treatment Definition: Eviction Notices

	(1)	(2)	(3)	(4)	(5)
Single Family	20.12***				
	(2.803)				
Old Condos		16.94***			
		(3.035)			
Mod. Condos			16.98***		
			(3.023)		
New Condos				19.96***	
				(2.917)	
New Construction					20.02***
					(2.927)
Constant	45.36***	45.36***	45.36***	45.36***	45.36***
	(3.818)	(4.160)	(4.114)	(3.849)	(3.868)
N	275	275	275	275	275
R^2	0.883	0.875	0.875	0.882	0.882
Treatment Mean	1799.1	2143.8	2319.7	1769.7	1713.4
Wild Bootstrap P-Value	0	0	0.00100	0	0

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates from difference-in-difference regressions on eviction notices using alternative measures of treatment to address concerns about mis-measurement. Column 1 assumes all new single family homes build between 1995 and 1999 replaced a duplex. Column 2 assumes all condos build before 1980 were converted in the post period. Column 3 assumes all condo modifications were condo conversions that replaced rent controlled units. Column 4 assumes all new condos replaced rent controlled units. Column 5 assumes all new construction in the post period replaced rent controlled units.

Table A6. Robustness to Treatment Definition: Wrongful Eviction Claims

	(1)	(2)	(3)	(4)	(5)
Single Family	7.684***				
	(2.299)				
Old Condos		6.302***			
		(2.141)			
Mod. Condos			6.380***		
			(2.115)		
New Condos				7.537***	
				(2.345)	
New Construction					7.606***
					(2.341)
Constant	11.56***	11.56***	11.56***	11.56***	11.56***
	(1.873)	(1.937)	(1.943)	(1.894)	(1.884)
N	275	275	275	275	275
R^2	0.831	0.821	0.822	0.828	0.829
Treatment Mean	1799.1	2143.8	2319.7	1769.7	1713.4
Wild Bootstrap P-Value	0	0.00100	0	0	0

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates from difference-in-difference regressions on eviction notices using alternative measures of treatment to address concerns about mis-measurement. Column 1 assumes all new single family homes build between 1995 and 1999 replaced a duplex. Column 2 assumes all condos build before 1980 were converted in the post period. Column 3 assumes all condo modifications were condo conversions that replaced rent controlled units. Column 4 assumes all new condos replaced rent controlled units. Column 5 assumes all new construction in the post period replaced rent controlled units.