American Economic Journal: Economic Policy 2018, 10(2): 62–89
https://doi.org/10.1257/pol.20150176

How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality?†

By Robert Collinson and Peter Ganong*

US housing voucher holders pay their landlord a fraction of household income and the government pays the rest, up to a rent ceiling. We study how two types of changes to the rent ceiling affect landlords and tenants. A policy that makes vouchers more generous across a metro area benefits landlords through increased rents, with minimal impact on neighborhood and unit quality. A second policy that indexes rent ceilings to neighborhood rents leads voucher holders to move into higher quality neighborhoods with lower crime, poverty, and unemployment. (*JEL* I38, R23, R31, R38)

A central goal of US low-income housing programs in recent years has been to improve neighborhood quality for assisted households. Recent evidence suggests this is a valuable goal, finding that neighborhood quality during childhood plays a role in determining labor market success as an adult (Chetty, Hendren, and Katz 2016; Chetty and Hendren 2016; Chyn 2016). The Housing Choice Voucher program tries to achieve this aim by providing households with more choice over location (US Department of Housing and Urban Development 2014). However, most housing voucher holders opt to live in neighborhoods of much lower quality than the average neighborhood, and typically live in neighborhoods similar to their neighborhood before receiving a voucher.[1] Various reforms to the generosity of vouchers have been proposed to address this problem, but little is known about whether these reforms achieve their goal of improving voucher holder neighborhood quality or are instead captured by landlords via higher rental prices.

*Collinson: Wagner School of Public Service, New York University, 295 Lafayette Street, New York, NY 10012 (email: rcollinson@nyu.edu); Ganong: Harris School of Public Policy, 1155 E 60th Street, Chicago, IL 60637 (email: ganong@uchicago.edu). This paper was previously circulated as "The Incidence of Housing Voucher Generosity." We thank Geoff Newton, Todd Richardson, Lynn Rodgers, and Lydia Taghavi at HUD, and MaryAnn Russ and Matt Hogan at DHA for answering many questions. We thank Nathaniel Baum-Snow, Raphael Bostic, Raj Chetty, Denise Dipasquale, Ingrid Gould Ellen, Michael Eriksen, Dan Fetter, Edward Glaeser, Colin Gray, Adam Guren, Nathan Hendren, Annie Levenson, Jeff Liebman, Jens Ludwig, Bruce Meyer, Pascal Noel, Larry Katz, Pat Kline, Ed Olsen, Jesse Rothstein, Barbara Sard, Dan Shoag, Neil Thakral, and seminar participants for their valuable feedback. Peter Ganong gratefully acknowledges residence at MDRC when working on this project, as well as funding from the NBER Pre-Doctoral Fellowship in Aging and Health, the Joint Center for Housing Studies, and the Taubman Center on State and Local Government. The views expressed here are those of the authors and should not be construed as representing those of the US Department of Housing and Urban Development or MDRC.

†Go to https://doi.org/10.1257/pol.20150176 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] We use "low quality" to refer to neighborhoods with low rents, high poverty, high crime, and poor performing schools.

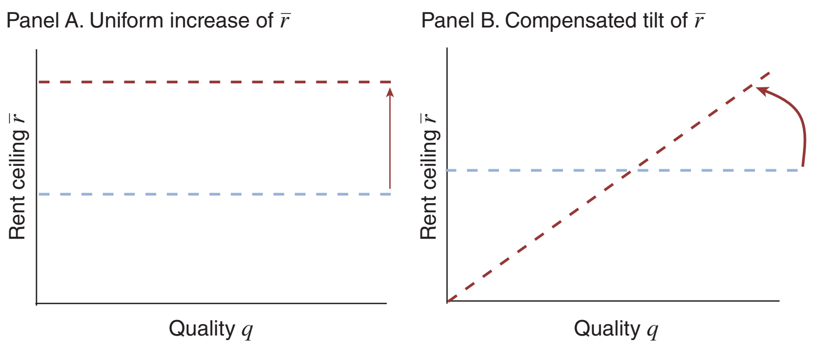

Figure 1. Changes in Rent Ceiling

Notes: This figure shows the two changes in voucher generosity that we study in this paper. The first increases the maximum per-unit government subsidy—which we refer to as the "rent ceiling"—uniformly in all neighborhoods in a metro area. The second increases the ceiling in higher quality ZIP codes and lowers it in lower quality ZIP codes.

We fill this void by evaluating two types of policy changes intended to spur moves to high-quality neighborhoods. The first increases the maximum per unit government subsidy, which we refer to as the "rent ceiling," uniformly in all neighborhoods in a metro area. The second increases the ceiling in higher quality ZIP codes and lowers it in lower quality ZIP codes. Each of these policy changes is depicted visually in Figure 1. We find that a policy of uniform increases in the ceiling raises the rents charged by voucher landlords to the government, with little impact on observed neighborhood quality. In contrast, a policy that establishes ZIP code-specific ceilings leads landlords to adjust rents, but is also a cost-effective way to increase neighborhood quality for voucher holders.

Housing Choice Vouchers, also known as Section 8 vouchers, paid rent subsidies for 2.3 million low-income families in 2016. Voucher holders typically pay 30 percent of their income as rent and the government pays the rest up to a rent ceiling, which is usually set at the fortieth percentile of metro area or countywide rents. Because a single uniform ceiling often applies to a broad geography, a much larger share of units are affordable with a voucher in low-quality neighborhoods. In 2013, census rent data show that two-thirds of rental units were priced at or below the ceiling in low-quality neighborhoods, but only one-seventh of units were in high-quality neighborhoods, as shown in Figure 2.

In spite of the importance of high-quality neighborhoods for economic mobility, most voucher households occupy units in low-quality neighborhoods. For example, we document that voucher holders in Dallas live on average in neighborhoods that are 1 standard deviation below the mean in terms of a neighborhood quality index defined below. Other research has shown that housing vouchers do not lead households to move to substantially safer or less impoverished neighborhoods.[2]

[2] Two examples with random assignment of housing vouchers are a lottery in Chicago (Jacob, Ludwig, and Miller 2013) and HUD's Welfare to Work Voucher Experiment (Eriksen and Ross 2013, Patterson et al. 2004). Two other studies that use matching methods are Carlson et al. (2012) and Susin (2002).

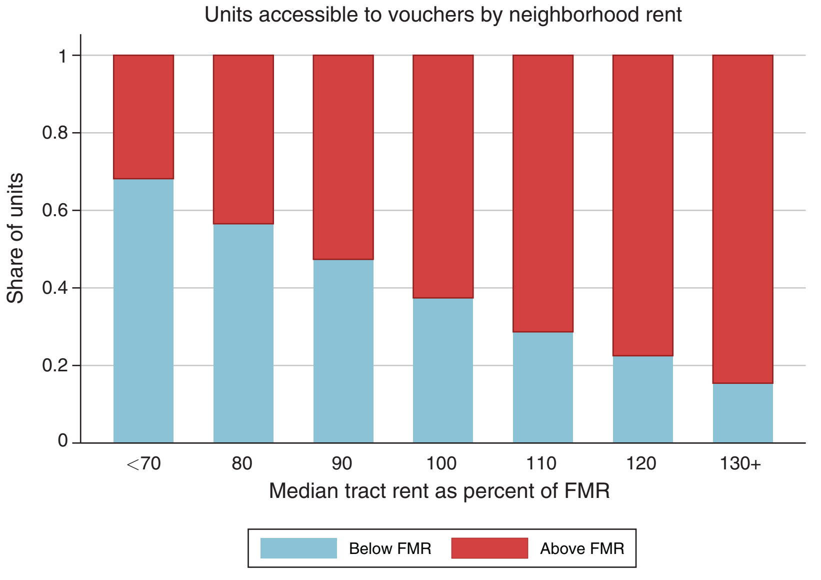

Figure 2. Unit Availability and Rent Distribution

Notes: Each year, the federal government publishes "Fair Market Rents." These are typically estimated as the fortieth percentile of rent in a county for studios, 1 bedroom, 2 bedroom, 3 bedroom, and 4 bedroom units. For each census tract, we compute the share of rental units priced at or below the fortieth percentile of the metro area rent distribution. This figure shows the average fraction of units priced below the rent ceiling as a function of median tract rent. Data are drawn from a special tabulation of the 2009–2013 ACS five-year estimate and FY2013 Fair Market Rents.

Who benefits from raising the rent ceiling uniformly is ambiguous.[3] It could benefit landlords, if they price discriminate by raising their rents to the new rent ceiling, or benefit voucher holders, if they use the more generous vouchers to move to better neighborhoods. Whether voucher holders move depends on the extent to which they value finding a unit in a high-quality neighborhood versus finding a unit at all. Ultimately, this is an empirical question, which we analyze using rich administrative and survey data.

In contrast to a uniform increase, tilting the rent ceiling so that it is higher in high-quality neighborhoods and lower in low-quality neighborhoods may be a cost-effective way to raise neighborhood quality.[4] Intuitively, the status quo penalizes searching in high-quality neighborhoods and the tilting policy raises optimal neighborhood quality by reducing this penalty. However, because the average rent ceiling does not increase, the scope for additional price discrimination is limited. In our empirical work, we investigate whether these two predictions are supported in the data.

[3] In online Appendix A, we theoretically analyze the impact of this policy using a model in which voucher holders face a trade-off between finding a unit in a high-quality neighborhood and finding a unit at all, and landlords can post higher rents in hopes of leasing to price-insensitive voucher holders. In such a model, whether landlords or tenants benefit more is ambiguous.

[4] The model in online Appendix A predicts that tilting the rent ceiling is a cost-effective way to raise neighborhood quality.

We empirically estimate the impact of the two voucher policies described above: raising the rent ceiling uniformly and tilting the rent ceiling towards quality. To estimate the impact of uniform increases, we use two complementary research designs; the first precisely measures the policy's impact on neighborhood quality, while the second uses a dataset with rich measures of unit quality. The first research design uses sharp corrections to accumulated measurement error in the local rent ceiling and a national panel capturing the universe of voucher holders. We estimate that a \$1 increase in the rent ceiling raises rents by \$0.46 over the next 6 years, while a hedonic measure of unit and neighborhood quality rises by only \$0.05 over the same time period. In addition, we estimate a precise zero for the impact on neighborhood quality as measured by census tract median rent and tract poverty rate. These point estimates imply that the benefit of this policy to landlords is eight times as large as the benefit in terms of observed quality to tenants. Although this design has the advantage of generating statistically precise estimates of the impact on neighborhood quality in an event study framework, it uses unit quality measures that are quite limited.

The second research design for studying a uniform rent ceiling increase remedies the limited unit quality measures in the first by exploiting a unique survey of voucher recipients. This survey of over 300,000 voucher holders has excellent detail on unit quality, including 26 questions on time-varying unit quality. We use a difference-in-differences design to study how unit quality changes in 39 metro areas that saw an increase in rent ceilings. Here, we find that each \$1 increase in the rent ceiling raised the rents paid on voucher units by \$0.47, with no significant impact on observed unit quality. These point estimates are very close to the point estimates from the first research design, although the estimates from the second research design are less precise. Two distinct research designs in two different time periods yield similar results: uniform increases in the rent ceiling appear to benefit landlords and not tenants.[5]

Finally, we study the effects of tilting the rent ceiling by examining a recent demonstration project in the Dallas, Texas metro area. Housing authorities in Dallas switched from a single metro-wide ceiling to ZIP-code-level ceilings in 2011. Much as with the uniform rent ceiling increase, we find empirically that landlords adjust rents—raising them in expensive ZIP codes and lowering them in low-cost ZIP codes. Because this policy makes vouchers more generous when they are used in high-quality neighborhoods, one might expect that it would improve neighborhood quality.

A difference-in-differences design using neighboring Fort Worth, Texas as a comparison group shows that new leases signed after the policy was implemented were in tracts where neighborhood quality is 0.23 standard deviations higher than leases signed prior to policy implementation. We construct a neighborhood quality index using the violent crime rate, test scores, the poverty rate, the unemployment rate, and the share of children living with single mothers. Relative to other housing voucher policies, 0.23 standard deviations is a substantial improvement in neighborhood quality. It is about half the magnitude of the improvements in neighborhood quality for people currently living in public housing who are allocated vouchers (Kling, Ludwig,

[5] These research designs estimate who benefits from *marginal* changes to the rent ceiling. See Desmond and Perkins (2016) for estimates of differences in *average* rents between similar voucher and non-voucher units in Milwaukee.

and Katz 2005) and larger than the improvement in neighborhood quality from allocating a voucher to previously unsubsidized tenants (Jacob and Ludwig 2012).

The Dallas tilting policy is budget-neutral within the time period we study. Absent any tenant behavioral response, this policy would have been cost-saving for the government because voucher holders tend to live in inexpensive neighborhoods, and therefore rent increases in expensive ZIP codes were offset by larger decreases in low-cost ZIP codes. Incorporating tenants' improved neighborhood choices, the Dallas intervention had zero net cost to the government over the years that we study. Thus, our results show that a simple budget-neutral reform to housing voucher design has the potential to substantially improve voucher holder neighborhood quality.

The remainder of the paper is organized as follows. Section I reviews the voucher program and Section II describes the data. In Section III, we show that a uniform increase in rent ceilings fails to raise neighborhood quality, but benefits landlords through increased voucher rents. In Section IV, we show that tilting rent ceilings is successful at inducing moves to higher quality neighborhoods. Section V concludes.

I. Housing Voucher Program

Housing Choice Vouchers use the private market to provide rental units for 2.3 million low-income households. There are four key actors in the voucher program: the US Department of Housing and Urban Development (HUD), local housing authorities, private landlords, and tenants. HUD funds local housing authorities that administer the voucher program, which includes making payments on behalf of tenants to landlords. Tenants search for units to lease on the private market.

The tenant pays at least 30 percent of her income in rent and the housing authority pays the difference, up to a rent ceiling. The local housing authority chooses a Payment Standard (which we refer to as the "rent ceiling") from 90 percent–110 percent of a federally-set "Fair Market Rent" (FMR) (Quadel Consulting Corporation 2001). HUD typically sets FMRs at the fortieth percentile of area-by-bedroom level gross rent (rent to landlord plus utility costs). By default, an FMR area is defined using county boundaries, but in urban areas there is often a single FMR for all counties in a metro area. We defer a discussion of how FMRs are updated until Section III, where we describe the natural experiments that we exploit.

Voucher holders renting units below the rent ceiling generally pay nothing when rents rise; the housing authority pays each extra dollar of such a rent increase. This is important because when the rent ceiling rises landlords can increase rents without worrying that this will cause the voucher holder to move. Two institutional details limit the extent of rent increases when the rent ceiling rises. First, a small share of voucher holders lease units with rents above the rent ceiling, and they bear each dollar of a rent increase. Second, at initial lease signing, as well as with requests for rent increases, housing authority staff must certify that rent requests meet "rent reasonableness" standards.[6]

[6] The typical rent reasonableness process entails local housing authority staff drawing a set of rent comparables for the unit in question from rental listing services. The housing authority staff will negotiate with a landlord requesting a rent substantially above the comparables, and may request evidence of other existing leases to establish

II. Data

The primary dataset we use in this paper is a HUD internal administrative database called "PIH Information Center" (PIC) that covers the universe of voucher holders. It contains an anonymous household identifier, an anonymous address identifier, building covariates, the rent ceiling, the FMR, and the contract rent received by a landlord on an annual basis, beginning in 2002. The data have two strengths that we exploit in our analysis. First, we can follow a household if they move in response to a policy change. Second, the address identifier, coded as a nine-digit ZIP code, enables us to follow a single address over time if it has multiple voucher occupants, which is useful for estimating the impact of an increase in the rent ceiling while holding constant many aspects of unit quality. Table 1 provides summary statistics and online Appendix B.1 discusses sample construction.

We supplement PIC with four other datasets. To investigate the effects of rent ceiling changes on non-voucher rents, we draw on rent data from the American Community Survey (ACS). To measure housing quality, we compute hedonic quality measures using coefficients from hedonic regressions in the ACS (Section IIIA) and American Housing Survey (Section IIIB). Our analysis in Section IIIB uses the predecessor to PIC, the Multifamily Tenant Characteristics System (MTCS), which contains information on voucher rents, location of voucher tenants, household size, and bedroom count. It also uses the HUD Customer Satisfaction Survey (CSS), which includes detailed questions about housing unit quality ideally suited to measure within-unit quality changes. To evaluate the effects of tilting the rent ceiling in Section IV, we assemble detailed data on neighborhood quality: school-level test scores data from the Department of Education, geocoded address level crime data from the Dallas Police Department, and tract-level measures from the American Community Survey 2006/2010.

III. Impact of Raising the Rent Ceiling Uniformly

We estimate the causal effect of uniform rent ceiling changes on neighborhood and unit housing quality and on voucher rents using two natural experiments. In Section IIIA, we study a 2005 change in FMRs due to availability of updated 2000 Decennial census data. The primary advantage of this research design is that it uses exogenous variation across all US counties, giving us enough statistical power to detect even small neighborhood quality responses. A secondary advantage is that by using unit fixed effects, we are able to examine the price response while holding physical structure quality and neighborhood quality constant. However, this design lacks detailed measures of within-unit quality changes arising from better management, maintenance, or unit upgrades.

that the requested rent is in line with market rents. The median housing authority rejects between one-quarter and one-half of units on the first inspection (Finkel and Buron 2001, Exhibit 3–5). One piece of evidence that the rent reasonableness process is effective is that empirically rents are lower for units with lower hedonic unit and neighborhood quality (online Appendix Figure B.1).

Table 1—Summary Statistics for Uniform Rent Ceiling Changes

	Mean (1)	SD (2)	Mean (3)	SD (4)
Rebenchmarking[a]				
Voucher characteristics	2004 (n = 1,578,124)		2010 (n = 1,665,868)	
Contract rent	495	238	586	266
Utility allowance	106	65	144	89
Rent ceiling (contract rent + utility)	618	278	762	296
Tenant payment	238	154	288	184
Tenant HH income (annual)	9,683	6,358	11,567	7,347
Share moved | nonattrit.	0.21	0.41	0.16	0.36
Tract characteristics[b]				
Poverty rate (2000)	16.31	9.13	16.02	9.07
Median contract rent (2005–2009)	473.70	196.26	479.55	197.97
Share voucher (2004)	0.021	0.024	0.019	0.022
County characteristics				
Fair market rent	628	312	802	326
Fortieth → fiftieth percentile FMRs[c]	Pre (n = 171,248)		Post (n = 285,279)	
Gross rent	547	167	620	213
Hedonic quality (using 28 survey vars)	613	237	628	247
Fair market rent	589	186	648	242

[a] Voucher and tract characteristics are computed giving equal weight to each county-bed pair.
[b] Poverty rate from the 2000 census, ACS survey responses from 2005 to 2009, with rent values inflated to 2009 dollars.
[c] "Pre" sample is 1999 and 2000. "Post" sample is 2001–2003. Summary statistics give equal weight to each county.

In Section IIIB, we investigate potential within-unit quality improvements. We make use of a detailed HUD survey that asked 26 questions about time-varying unit quality and was administered to voucher holders on a widespread basis from 2000 to 2003 to construct measures of housing quality. Here, we exploit a 2001 change that raised FMRs from the fortieth percentile to the fiftieth percentile of rents in 39 metro areas. Across both research designs, we find similar results: raising the rent ceiling results in higher rents with little evidence of positive unit or neighborhood quality impacts. We discuss at the end of the section why we believe that price discrimination by landlords is the most reasonable interpretation of these empirical results.

A. *Rebenchmarking of FMRs in 2005*

For many years, data constraints meant that FMRs changed little in a typical year, punctuated by very large swings once every ten years. This offers useful variation for a quasi-experimental analysis. In most years, FMRs are updated using local Consumer Price Index (CPI) rental measures for 26 large metro areas and 10 regional Random Digit Dialing (RDD) surveys for the rest of the country. The availability of new decennial census data results in a "rebenchmarking." Because the local CPI and RDD estimates are noisy, large swings in FMRs occurred from 1994 to 1996 when 1990 census data were incorporated into FMRs, and again in 2005 when 2000 census data were added. In non-rebenchmarking years, FMR changes are very crude estimates of the actual change in local rent; for example, they were

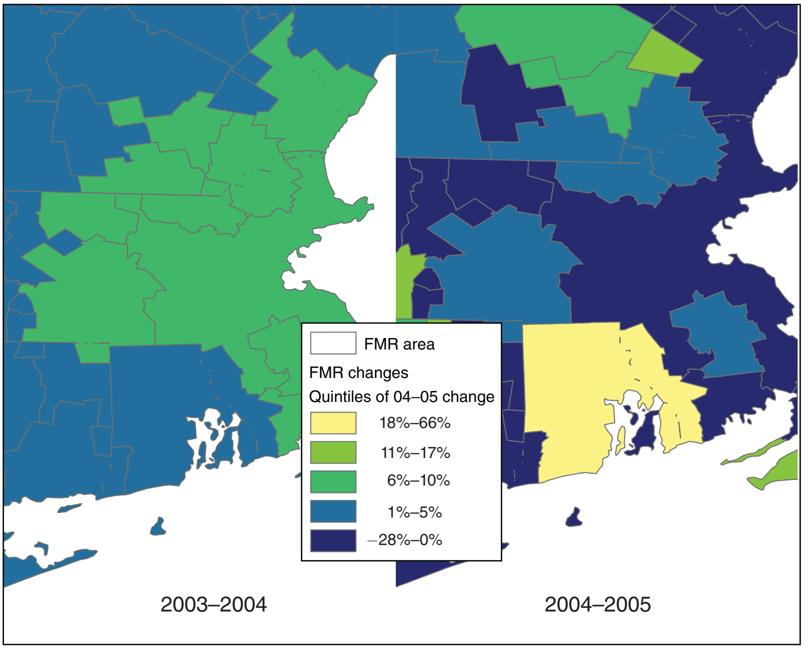

Figure 3. Eastern New England FMR Changes

Notes: This map shows changes in fair market rents from 2003 to 2004 and 2004 to 2005. In most years, including 2003 to 2004, one inflation factor is used for Greater Boston and another is used for all of eastern New England. In 2005, the government made large revisions as part of a "rebenchmarking" to incorporate newly available county-level data from the 2000 census.

a bit *worse* at predicting local rent changes than using a single national trend from 1997 to 2004.[7]

The 2005 rebenchmarking offers substantial variation in FMR changes, suitable for a quasi-experimental research design. As an example, in Figure 3, we show FMR revisions for two-bedroom units in Eastern New England for 2003–2004 and for 2004–2005. From 2003 to 2004, FMRs rose by 6 percent in Eastern Massachusetts and rose by 2 percent in outlying areas. The next year shows large revisions, with Rhode Island experiencing 22 percent increases in 2-bedroom FMRs and Greater Boston experiencing 11 percent decreases. Figure 4 shows national impacts of the rebenchmarking. Figure 5 shows an event study of FMRs for four groups of county-bed pairs, stratified by the size of their revision from 2004 to 2005. In nominal terms, the bottom quartile fell by 7 percent, while the top quartile rose by 24 percent. These four groups had similar trends in the six years after the revision, so we can study the rebenchmarking as a one-time, permanent change. Throughout the paper, all regression specifications studying rent or hedonic quality use a log transformation. The motivation for this log transformation is that there is

[7] The top of the panel in online Appendix Figure B.2 shows that the variance of FMR changes is much larger in rebenchmarking years. The bottom panel shows that using a single national trend instead of actual FMR changes would have resulted in smaller swings in rent in the 2005 rebenchmarking.

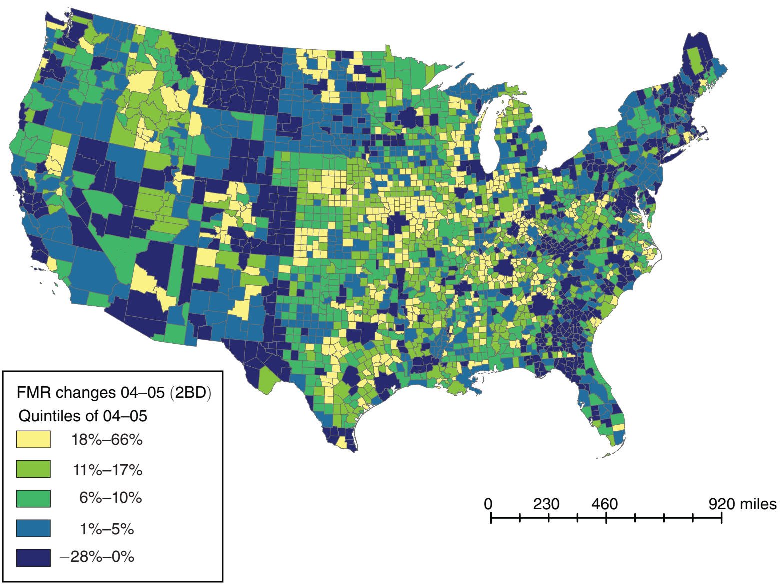

Figure 4. National Fair Market Rent Rebenchmarking, 2004–2005

Notes: This map shows changes in fair market rents from 2004 to 2005. In 2005, the government made large revisions as part of a "rebenchmarking" to incorporate newly available county-level data from the 2000 census.

tremendous heterogeneity in FMR levels; in 2004, FMR levels for a 2-bedroom unit ranged from \$370 in rural Alabama to \$1,800 in San Jose. A \$50 increase in the FMR would have a very different impact in percent terms in Alabama than in San Jose. Additional empirical details on our use of the rebenchmarking are provided in online Appendix B.2.

To clarify the sources of variation that we use for identification, we show that the rebenchmarking can be decomposed into three pieces: changes in non-voucher rents, measurement error from annual updates, and measurement error in the census. Define σ_t as an annual estimate of the change in log rents based on a regional RDD or CPI survey from year $t-1$ to t.[8] Define $\exp(r_t + \varphi_t)$ as an observation from decennial census data, where $\exp(r_t)$ is the true rent and $\exp(\varphi_t)$ is census measurement error. We can use these definitions to write $\log FMR^{2004} = \sum_{t=1991}^{2004} \sigma_t + r_{1990} + \varphi_{1990}$, and $\log FMR^{2005} = \sum_{t=2001}^{2005} \sigma_t + r_{2000} + \varphi_{2000}$. Taking the difference gives

$$\Delta FMR = \underbrace{r_{2000} - r_{1990}}_{\text{true rent change}} + \underbrace{\sigma_{2005} - \sum_{t=1990}^{1999} \sigma_t}_{\text{annual meas. error}} + \underbrace{(\varphi_{2000} - \varphi_{1990})}_{\text{census meas. error}}.$$

Consistent with measurement error as a source of variation, places where FMRs drifted upward due to noise over the prior ten years were subject to *downward*

[8] The RDD and CPI surveys are used to produce adjustment factors that modify the base, not to provide a new estimate of the level.

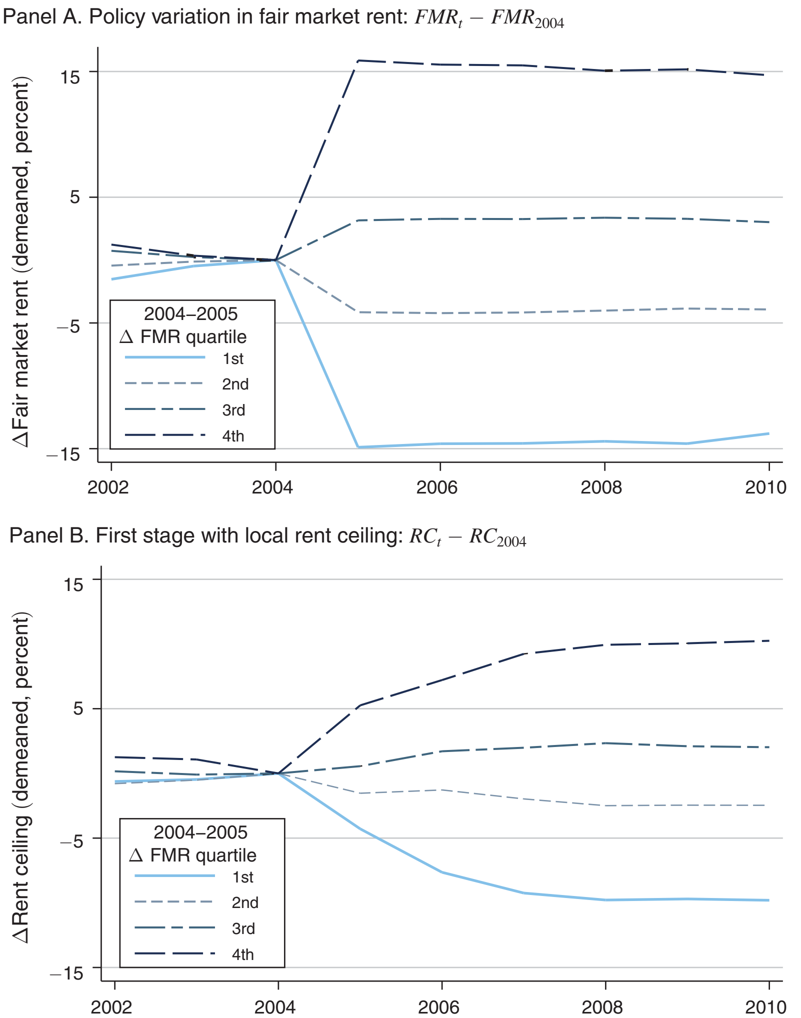

Figure 5. Event Study for Rebenchmarking

Notes: In 2005, the government made large revisions as part of a "rebenchmarking" to incorporate newly available data from the 2000 census. Panel A plots demeaned changes in the fair market rent for four quartiles of county-bed observations, stratified by the change from 2004 to 2005. Local housing authorities administer the vouchers, and have discretion to set the local rent ceiling at 90 percent, 100 percent, or 110 percent of fair market rent. Panel B plots local rent ceilings, using the same grouping of county-beds as in panel A.

revisions in 2005, and places where FMRs drifted downward due to noise were subject to *upward* revisions.

Suppose that outcomes y, such as unit and neighborhood quality or voucher rents, may be affected by the rent ceiling $\overline{r}$ as well as contemporaneous shocks to supply and demand η, as expressed by the empirical model $\Delta y = h(\overline{r}) - h(\overline{r}_{2004}) + \eta$. Our identifying assumption is the shocks *after* 2004 were orthogonal to the level of FMRs in 2005, conditional on their 2004 level.

ASSUMPTION: (Identification Assumption in Rebenchmarking Research Design):

$$\eta \perp FMR_{2005} \,|\, FMR_{2004}.$$

As detailed above, ΔFMR consists of measurement error, which is by construction orthogonal to future trends, and the true non-voucher rent change, $r_{2000} - r_{1990}$. Note that this research design allows the rebenchmarking to bring rental rents closer in line with the *level* of market fundamentals. We require only that the *change* in FMR be uncorrelated with the subsequent shocks η. Available empirical evidence supports this identification assumption. First, rents are about flat from 2002 to 2004, prior to the policy change. Second, contemporaneous changes in non-voucher rents have no significant correlation with the FMR change.[9]

We estimate an empirical specification using two-stage least squares because local housing authorities have some discretion in setting rent ceilings, as discussed in Section II. Formally, we estimate a first stage

$$\text{(1) First Stage:} \quad \overline{r}_j = \alpha + \gamma FMR_{2005j} + \omega FMR_{2004j} + \kappa \overline{r}_{2004j} + \varepsilon_j,$$

where we predict the rent ceiling for county-bed j with the *FMR* for j in 2004 (FMR_{2004j}), the rent ceiling $\overline{r}$ for j in 2004, and exogenous variation from the 2005 FMR for j (FMR_{2005j}) with error term ε.[10] In the short term, housing authorities use their discretion in setting rent ceilings to offset the immediate impact of FMR changes, but a \$1 increase in the FMR from 2004 to 2005 corresponded to a \$0.58 increase in the rent ceiling by 2010, as estimated by coefficient $\hat{\gamma}$. We estimate our second stage where j indexes county-bed pairs, $\hat{\overline{r}}_j$ is the fitted value from the first stage equation, and the coefficient of interest is β the effect of rent ceiling changes on the outcome Δy_j:

$$\text{(2) Second Stage:} \quad \Delta y_j = \alpha + \beta \hat{\overline{r}}_j + \lambda FMR_{2004j} + \pi \overline{r}_{2004j} + \eta_j.$$

We assess the effects of uniform rent ceiling changes on neighborhood quality as measured by median tract rent, neighborhood quality as measured by tract poverty rate, rents received by landlords, and a "composite" hedonic measure of unit and neighborhood quality.[11] Tract-level measures are a good way to detect even small improvements in neighborhood quality because census tracts typically have 4,000

[9] Online Appendix B.3 analyzes prior and contemporaneous changes in non-voucher rents in more detail and online Appendix Table 1 shows the relevant regression results.

[10] The motivation for controlling for 2004 FMR is driven by the nature of our quasi-experimental variation. Prior to the FMR change, average rents across all units were *rising* for places about to receive a downward revision and that rents were *falling* for places about to be revised upward; this was likely because of mean reversion in regional rents combined with infrequent FMR resets. Controlling for the 2004 FMR level eliminates this pretrend. We also try the following first-differences specification. We estimate a first stage: $\Delta \overline{r}_j = \alpha + \gamma \Delta FMR_j + \varepsilon_j$, where $\Delta \overline{r}_j = \overline{r}_j - \overline{r}_{2004j}$, and second stage: $\Delta y_j = \alpha + \beta \widehat{\Delta \overline{r}_j} + \eta_j$. This specification produces very similar point estimates.

[11] We use the term "composite" hedonic quality when the measure incorporates characteristics of both the unit (such as building age and type) and neighborhood (such as tract median rent).

residents and 77 percent of voucher moves cross tract boundaries.[12] We construct our voucher rent measure as $\Delta y_{t,j} = r_{t,j}^{voucher} - r_{2004,j}^{voucher}$.

To construct our composite hedonic quality measure, we run a hedonic regression in the American Community Survey using covariates for the number of bedrooms, structure age, structure type (e.g., single family, multi-family, or apartment building), and neighborhood rent.[13] We then construct our dependent variable quality measure $\Delta y_j = \hat{\beta}_{hedonic}(x_{t,j} - x_{2004,j})$ using covariates $x_{t,j}$ on the number of bedrooms, structure type, structure age, and median tract rent from the voucher data, where $x_{t,j}$ is the unconditional average of x in county-bed j in year t.[14]

The impact of raising the ceiling on observable quality is very small. Table 2, columns 1–3 show the effects of a \$1 change in the rent ceiling on neighborhood and unit quality. A \$1 increase in the ceiling has no economically significant impact on the neighborhood quality of voucher tenants, as measured by neighborhood rents (column 1) or poverty rates (column 3), and raises composite hedonic quality by a mere \$0.05 cents (column 2).

In contrast to the quality results, average rents rise by \$0.46 cents in response to a \$1 increase in the rent ceiling (Table 2, column 4). Figure 6 plots the year-by-year coefficients of the reduced-form impact of the FMR change on rents, and shows rents rise steadily in response to the rent ceiling increase through the first four years after the rebenchmarking, while composite hedonic quality rises minimally throughout this period. These results imply that only $(0.05/0.46 =)$ 11 percent of the increased government expenditure went to improvements in observable unit or neighborhood quality.

We conduct three robustness checks of our finding that landlords adjust rents in response to rent ceiling changes.[15] First, we address the concern that places revised upward might have different rent fundamentals than places revised downward. To do this we add county fixed effects to equations (1) and (2) so that identification comes only from within-county variation comparing the FMR change by bedroom count. Our point estimates from the model with county fixed effects of \$0.50 cents are remarkably similar to our baseline estimate of \$0.46.

[12] The tract rent measure is $\Delta y_{t,j} = \log(tract\, rent_{t,j}) - \log(tract\, rent_{2004,j})$, the difference in average median tract rent for vouchers in county-bed j from year 2004 to year t. The census tract poverty rate is $\Delta y_{t,j} = tract\, pov_{t,j} - tract\, pov_{2004,j}$, where $tract\, pov_{t,j}$ is the average tract poverty rate of voucher holders in county-bed j.

[13] Voucher holders are assigned an appropriate number of bedrooms according to a fixed schedule based on household size. We use this assigned bedroom count to construct our instrument values and in our county-bed definitions. A voucher holder can choose to lease a larger unit—for example, a family eligible for two bedrooms can lease a three bedroom unit—but the payment will be according to their eligible number of bedrooms. To capture moves to larger units, we include the actual number of bedrooms in the leased apartment as a quality measure.

[14] We estimate our hedonic coefficients in the American Community Survey, where the smallest geographic units are Public Use Microdata Areas (PUMAs) with about 150,000 residents. The results from our hedonic regression appear in online Appendix Table 2. When predicting composite hedonic quality for voucher units, we measure neighborhood quality using median tract rent. Substituting median tract rent for a PUMA fixed effect offers a much more granular neighborhood quality measure and likely has little impact on the other hedonic coefficients. To assess the potential change in hedonic coefficients from using median tract rent instead of a PUMA fixed effect, we re-run our hedonic regression using the median PUMA rent in lieu of the PUMA fixed effect. We find that the hedonic coefficients are largely unchanged, the coefficient on PUMA median rent is approximately \$1 and the constant shrinks from \$900 to \$50. More details on construction of the hedonic measure are provided in online Appendix B.4.

[15] The regression results are shown in online Appendix Table 3.

Table 2—Effect of Uniform Rent Ceiling Increase on Rent and Quality
(*research design: rebenchmarking*)

	Hedonic quality			
	Neighborhood rent (1)	Unit and neighborhood (2)	Neighborhood poverty (3)	Voucher rents (4)
	ΔY, 2004–2010			
	log median tract rent	log hedonic quality	Tract poverty	log voucher rent
IV estimate				
log rent ceiling 2010	0.029	0.047	−0.002	0.458
	(0.018)	(0.016)	(0.006)	(0.030)
	(5)			
First stage	Y: log rent ceiling 2010			
log FMR 2005	0.580			
	(0.037)			
mean(Y)	6.107	7.136	0.162	6.130
Unit of observation	County-bed	County-bed	County-bed	County-bed
Observations	12,333	12,333	12,333	12,333

Notes: This table shows the quality and rent impacts of a uniform increase in the rent ceiling using variation from the 2005 FMR rebenchmarking. Columns 1–4 report the results of estimating equation (2) from Section III on different dependent variables. Columns 1–3 report the effects of rent ceiling changes on changes to three housing quality measures for all voucher holders from 2004–2010. Hedonic quality in column 2 is based on structure age, structure type, number of bedrooms, and median tract rent (see Section VA for details). Column 4 reports the effect of the rent ceiling change on changes in voucher rents from 2004–2010. Column 5 reports the first stage from estimating equation (1) in Section III. The sample consists of all tenants where the unit of observation is county-bed pairs. Standard errors are clustered at the FMR group level.

Second, we show that it is the government and not voucher holders who pay more when the rent ceiling rises. Recall from Section I that some voucher holders choose to rent a unit that costs more than the rent ceiling and then pay more than 30 percent of their income. In this case, when the landlord raises rents, it is the voucher holder and not the government that pays an additional dollar, potentially undermining the interpretation that landlords are price discriminating on the basis of voucher receipt. We address this concern by building a sample of tenants that are unlikely to be the residual payer.[16] For this subsample, we can be confident that when rent rises by $1 that the government pays $1 more. A $1 increase in the rent ceiling raises rents by a similar amount to our baseline specification.

Finally, to address concerns about whether rent increases may reflect quality improvements not captured by our hedonic measure, we estimate a model with address fixed effects. The sample consists of 800,000 units continuously occupied by a voucher tenant (either a new voucher tenant or an existing tenant). Here, we find rent increases of $0.15 cents for each dollar increase in the rent ceiling. The

[16]To identify households that are unlikely to be the residual payer, we examine two variables: the gap between gross rents and the rent ceiling, and the number of bedrooms in 2004. We use voucher holders with two or fewer bedrooms and a value of rent minus rent ceiling in the bottom three quintiles in 2004. The probability that these households have rent higher than the rent ceiling—and therefore pay more when the landlord raises the rent—is 11 percent.

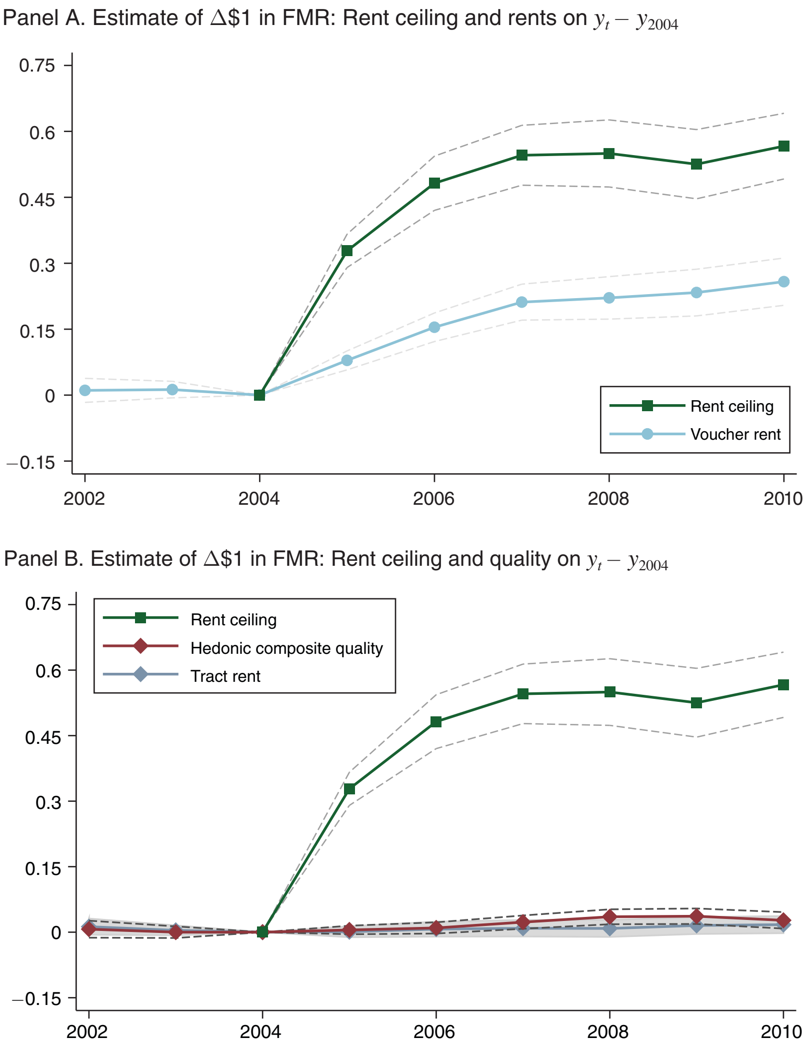

FIGURE 6. IMPACTS OF REBENCHMARKING: RENT AND QUALITY

Notes: We plot β coefficients from a reduced-form regression for rent ceilings, rents, and quality using the following equation $\Delta y_{t,j} = \alpha + \beta FMR_{2005,j} + \phi FMR_{2004,j} + v\bar{r}_{2004,j} + \varepsilon_j$. The top panel shows impacts on the rent ceiling compared to voucher quality. The bottom panel repeats the rent ceiling estimates from the top panel and also shows impacts on quality. The coefficient β captures the impact of a $1 increase in the FMR on each variable. Hedonic quality is measured using number of bedrooms, structure type, structure age, and median tract rent. Shaded area/dashed lines indicate 95 percent confidence intervals. Rental data from 2002 and 2003 are a test for pre-trends, and the 2004–2005 first stage is used. See Section IIIA for details.

address fixed effects specification indicates that rents increase when the rent ceiling rises, even after holding constant neighborhood quality and permanent unit attributes. There are two potential reasons why the address fixed effects estimate ($0.15) is smaller than the full sample estimate ($0.46). One explanation is that the government is more easily able to enforce the "rent reasonableness" restrictions discussed in Section I when the same unit was previously leased to a voucher recipient and so the government has an easily-available benchmark for what the unit's rent should be. A second explanation, which we investigate in detail in the next section and do not find any evidence for, is that increased rents pay for improvements in time-varying unit quality.

B. *Fortieth → Fiftieth Percentile FMRs in 2001*

A concern with the first research design is an inability to measure detailed elements of quality that might vary over time within the unit. In a different dataset, HUD measured unit quality in much more detail from 2000 to 2003. Using this dataset requires a different identification strategy based on a policy change in 2001, when HUD switched from setting FMRs at the fortieth percentile of the local non-voucher rent distribution to the fiftieth percentile in 39 metro areas. This policy was implemented not in response to recent housing market conditions, but rather with the explicit goal of "deconcentration" of vouchers from the lowest quality neighborhoods.[17]

From 2000 to 2003, HUD conducted a Customer Satisfaction Survey (CSS) of about 100,000 voucher households each year. This survey included numerous questions on unit quality and came close to matching the level of detail in the American Housing Survey (AHS), which is the state-of-the-art data source on housing quality in the United States. In particular, it asked many questions about unit attributes that could plausibly vary at the same address over time including: "How would you rate your satisfaction with your unit?," "Has your heat broken down for more than 6 hours?," "Does your unit have mildew, mold, or water damage?," and "Have you spotted cockroaches in your home in the last week?" A full list of quality measures is in online Appendix B.4. We transform these questions into a hedonic unit quality measure and a composite (unit and neighborhood) hedonic quality measure that includes tract median rents from the 2000 census. Our analysis pools these county-year observations from 1999–2003. To compute hedonic quality, we identified the 26 questions on time-varying quality in the CSS that also appeared in the AHS.[18] We run a hedonic regression in the AHS using these 26 questions, number of bedrooms, building age, and building type, and a measure of median neighborhood rent, and then use tenants' responses in the CSS to predict composite hedonic quality. We also assess the effects of the intervention on voucher rents using administrative records from PIC and its predecessor, the MTCS.[19] To construct our rents measure we calculate the average by county-year for all tenants.

[17] The 39 metro areas were chosen on the basis of three factors, which are not obviously related to the *trend* in voucher rents or neighborhood quality:

- a size requirement (must contain at least 100 census tracts);
- an FMR neighborhood access measure—70 percent or fewer of census tracts (with at least 10 two bedroom rental units) having at least 30 percent of the two bedroom rental units with gross rents at or below the two bedroom FMR; and
- a high concentration of voucher holders in a limited number of census tracts—25 percent or more of tenant-based voucher holders reside in 5 percent of tracts with FMR area with largest number of participants.

[18] Online Appendix Table 4 compares the predictive performance of our hedonic characteristics across datasets. In the AHS, the CSS variables perform nearly as well as the "kitchen sink" AHS model (R^2 0.31 for CSS variables compared to 0.42 for the full AHS model). See US Department of Housing (2000) for more details.

[19] We use the administrative data on rents because they cover the universe of voucher tenants. The CSS contains rents for survey respondents but the values are top-coded at $500 and reported in bands of $100.

We estimate the impacts of the fortieth to fiftieth percentile policy change on Fair Market Rents, actual voucher rents, and composite quality. In order to assess the impact of the rent ceiling increase, we implement a difference-in-differences model using an instrumental variable specification.[20] Our estimates of the policy's effects on housing quality use individual-level survey data from the CSS, and the effects on rents use administrative data aggregated to the county-level. In our first stage in equation (3), we predict the endogenous rent ceiling for household i in FMR area j and time $t\,(\overline{r}_{ijt})$ using an indicator for being in an FMR area subject to the fiftieth percentile FMR policy $(\mathbf{1}(FMR = 50_j))$, an indicator for whether time period t is after the policy change $(\mathbf{1}(Post_t))$, and the excluded instrument: an indicator for the whether the observation is in FMR Area subject to fiftieth percentile FMR after the policy change $(\mathbf{1}(FMR = 50_j \times Post_t))$.[21] Our second-stage question is represented by equation (4), where $\widehat{\overline{r}_{ijt}}$ is the fitted value from the first-stage (the predicted payment standard) and β is the parameter of interest, the effect of a policy-induced change in the rent ceiling on the outcome:

$$\text{(3) First Stage:} \qquad \overline{r}_{ijt} = \pi + \gamma\mathbf{1}(FMR = 50_j \times Post_t) + \mathbf{1}(FMR = 50_j) + \mathbf{1}(Post_t) + \varepsilon_{ijt};$$

$$\text{(4) Second Stage:} \qquad y_{ijt} = \alpha + \beta\widehat{\overline{r}_{ijt}} + \mathbf{1}(FMR = 50_j) + \mathbf{1}(Post_t) + \eta_{ijt}.$$

Our identification condition is the standard difference-in-differences condition: $E(\eta_{ijt}\,|\,\mathbf{1}(FMR = 50 \times Post)) = 0$. Figure 7 shows the results visually and Table 3 shows regression results. Setting FMRs at the fiftieth percentile of the local non-voucher rent distribution raised rent ceilings by an average of 11 percent. For every \$1 increase in FMRs, rents rose by \$0.47 (column 5), which is very similar to our estimate of \$0.46 when using the rebenchmarking research design. In comparison, composite hedonic quality rose by \$0.04 (Table 3, column 3), with a standard error of \$0.09. Although the estimate for the impact on quality is less precise than in the rebenchmarking research design, the results from this analysis reinforce the conclusions from the prior section that uniform rent ceiling increases in FMRs do not seem to improve quality.

Our empirical results from two separate natural experiments show that uniform changes in the ceiling do little to improve either neighborhood or observed unit quality for voucher tenants while increasing rents substantially. We interpret our findings as likely reflecting landlords price discriminating by raising rents in response to rent ceiling changes. Our empirical findings are also consistent with landlords improving unmeasured aspects of unit quality and raising rents to cover the cost of these

[20] A difference-in-difference specification estimating the average effect of the policy δ using the following equation, $y_{ijt} = \alpha + \delta\mathbf{1}(FMR = 50_j \times Post_t) + \mathbf{1}(FMR = 50)_j + \mathbf{1}(Post_t) + \eta_{ijt}$, appears in online Appendix Table 5.

[21] In the case where the outcome is the voucher rent our regressions are at the county-year level: $\overline{r}_{jt} = \pi + \gamma\mathbf{1}(FMR = 50_j \times Post_t) + \mathbf{1}(FMR = 50_j) + \mathbf{1}(Post_t) + \varepsilon_{jt}$, where j now indexes counties. Again, $\mathbf{1}(\cdot)$ denotes the indicator function, taking the value equal to 1 if the statement is true and zero otherwise.

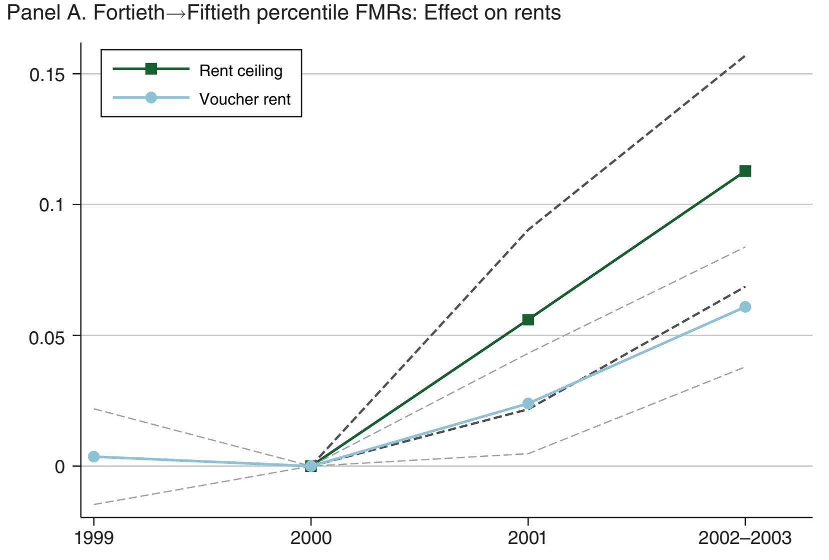

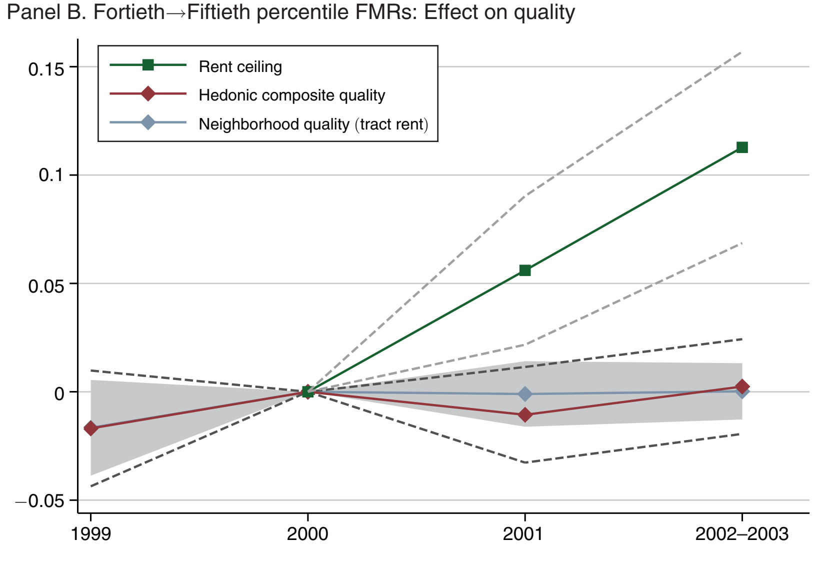

Figure 7. Impacts of Fortieth → Fiftieth Percentile FMRs: Rent and Quality

Notes: Panel A shows an event study for changes in the rent ceiling and voucher rents around the introduction of fiftieth percentile FMRs in 2001. Panel B plots the same event study for changes in quality: hedonic composite quality and neighborhood quality. Hedonic composite quality is measured using number of bedrooms, structure type, structure age, median tract rent, and 26 survey questions about unit quality and maintenance. Neighborhood quality is measured using median tract rents. Shaded area/dashed lines indicate 95 percent confidence intervals. See notes to Table 3 for details.

improvements. However, we view unmeasured quality improvements as unlikely to fully explain the estimated rent increases because we have very detailed measures of unit quality, and if a landlord decides to make unit improvements, then at least some of them would show up in the observable dimensions of unit quality.

Table 3—Effect of Uniform Rent Ceiling Increase on Rent and Quality (IV Estimate)
(*research design: fortieth → fiftieth percentile FMRs*)

	Hedonic quality				
	Neighborhood	Unit	Unit and neighborhood	Neighborhood poverty	Voucher rents
	(1)	(2)	(3)	(4)	(5)
	Y: log median tract rent	*Y*: log unit hedonic quality[a]	*Y*: log composite hedonic quality	*Y*: Tract poverty rate	*Y*: log rent[b]
log rent ceiling	0.054	0.005	0.041	−0.006	0.467
	(0.056)	(0.053)	(0.090)	(0.016)	(0.106)
Unit of observation	Household	Household	Household	Household	County
Observations	315,629	315,629	315,629	315,629	11,829

Notes: This table shows the quality and rent impacts of a metro-wide increase in the rent ceiling using variation from the fortieth → fiftieth percentile FMR change from 2000 to 2003. The sample is voucher households in the customer satisfaction survey in years 2000–2003 for columns 1–4. The sample for column 5 is all county-years with valid rent data in our pooled MTCS and PIC datasets. The table reports the effect of a $1 increase in the rent ceiling. Standard errors are clustered at the FMR group level. See Section III for details.

[a] Uses only the structural and time-varying components of quality, excludes neighborhood rent.

[b] Uses county-level average voucher rents from HUD's PIC and MTCS administrative datasets for 2000–2003.

IV. Tilting the Rent Ceiling with ZIP-Level FMRs in Dallas

In contrast to the results in the previous section, we find that tilting the rent ceiling has a big impact on neighborhood quality. Following a court settlement, HUD replaced a single metro-wide FMR in Dallas with ZIP code-level FMRs in early 2011. The new ZIP code-level FMRs were set by multiplying the metro-wide FMR in Dallas by the ratio of the median gross rent of rental units in the ZIP code to median gross rent of units in the metro area. The demonstration caused sharp changes in local rent ceilings, ranging from a decrease of 20 percent to an increase of 30 percent, as shown in the top panel of Figure 8.

In Section IVA, we validate that landlords in Dallas behave similarly to landlords nationally in response to uniform increases: voucher rents rose in ZIP codes where FMRs rose and fell in ZIP codes where FMRs fell. In Section IVB, we build a neighborhood quality index and show that households who moved located in neighborhoods 0.23 standard deviation higher in quality. Finally, in Section IVC, we compare the effects on neighborhood quality to the results from more costly alternative interventions. Online Appendix B.5 contains supplementary empirical details.

A. *Impacts on Voucher Rents and Building Quality*

We document that the ZIP-level elasticity of rents and building quality in response to changes in the rent ceiling in Dallas is similar to the responses to uniform rent ceiling increases. The rent results provide validation that landlords in Dallas respond similarly to landlords nationally when the rent ceiling changes. The identifying assumption for this analysis is that the relationship between the ZIP FMR and our

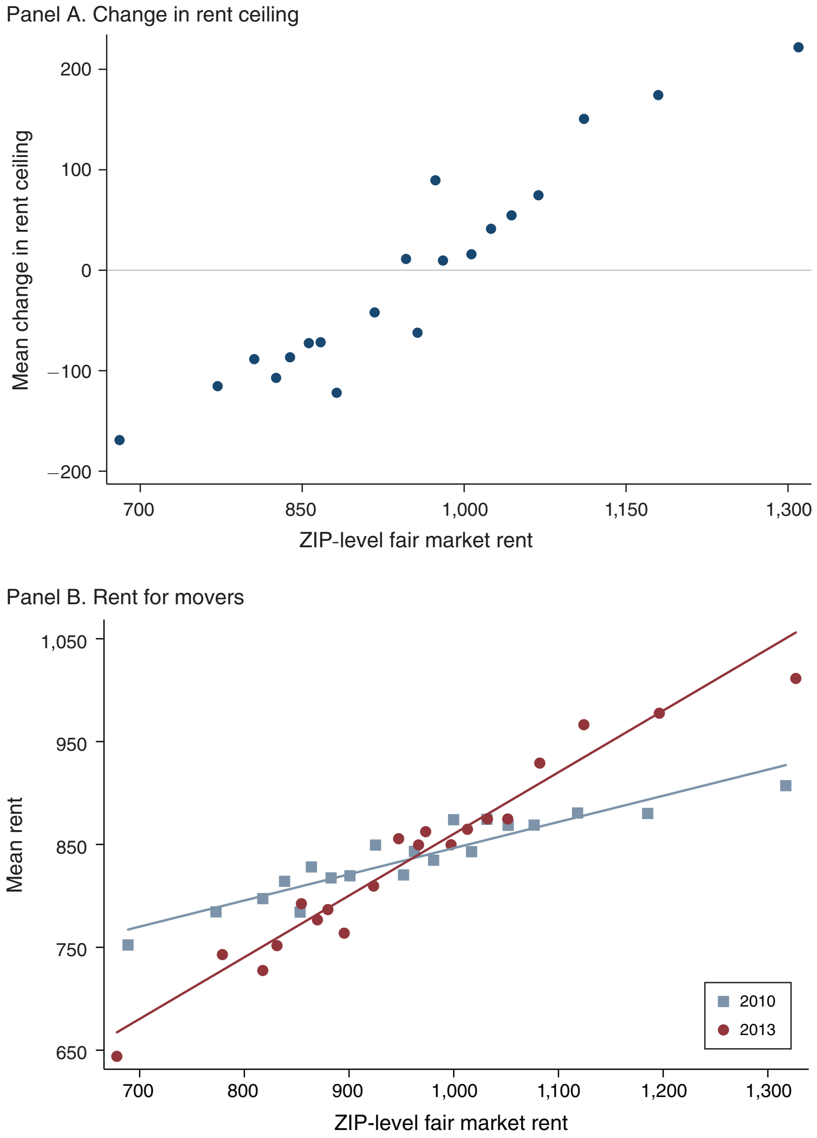

Figure 8. Impact of Tilting: Rent Ceiling and Rents

Notes: In 2011, Dallas replaced a single, metro-wide FMR with ZIP code-level FMRs. Panel A shows that this policy raised rent ceilings in expensive neighborhoods and lowered rent ceilings in cheap neighborhoods. Using a sample of households that moved from 2010 to 2013, we residualize ZIP FMRs and tenants' rent ceiling by the number of bedrooms, add back the unconditional mean for each, and plot conditional mean rent ceilings for 20 quantiles of residualized ZIP code-level FMR. Panel B plots mean rents against the ZIP-code level FMR for movers from 2010–2013 at their 2010 and 2013 ZIP codes. We follow the same procedures as above using residual voucher rents by bedroom size. Each dot reflects means for one of 20 quantiles of the ZIP code-level FMR distribution conditional on bedroom-year in 2010 and in 2013. Rents were quite responsive to the new rent ceiling schedule.

outcomes (housing quality and voucher rents) would be unchanged from the base year (2010) to the most recent data available (2013), but for the policy change:

ASSUMPTION: (Identification Assumption in ZIP Code-Level Research Design):

$$\eta \perp FMR \times Post | FMR.$$

Because FMR in 2010 was constant across Dallas, using the 2011 FMR level as the regressor is the same as using the change from 2010 to 2011 as the regressor. Our sample consists of voucher holders in 2010 and 2013. In our first stage we predict the payment standard for voucher holder i in ZIP code j at time t ($\overline{r}_{ijt}$) using equation (5). For voucher household i in ZIP code j in year $t \in \{2010, 2013\}$, $\mathbf{1}(Pos\ t_t)$ is a dummy for 2013, FMR_j is the applicable FMR level in 2011 for ZIP code j, and b_{ijt} is set of dummy variables for the number of bedrooms interacted with the year. The inclusion of this term eliminates the need for a separate year fixed effect term. We estimate

$$\text{(5) First Stage:} \qquad \overline{r}_{ijt} = \alpha + \gamma FMR_j \mathbf{1}(Pos\, t_t) + \omega FMR_j + b_{ijt} + \varepsilon_{ijt};$$

$$\text{(6) Second Stage:} \qquad y_{ijt} = \alpha + \beta \widehat{\overline{r}_{ijt}} + \lambda FMR_j + b_{ijt} + \eta_{ijt}.$$

Our second-stage equation (6) estimates the effect, β, of policy-induced changes in the payment standard on voucher rents or building quality (y_{ijt}). Rents at the ZIP code-level were highly responsive to the policy change, as shown in Figure 8. Online Appendix Table 6 reports results from equations (5) and (6). Changes in FMRs are a strong predictor of changes in rent ceiling, with a coefficient of \$0.62. We find substantial rent increases in more expensive areas and rent decreases in cheaper areas; every \$1 change in the rent ceiling caused a \$0.57 change in rents. This estimate is similar to the estimates in Section III that a \$1 change in the rent ceiling raised rents by \$0.46–\$0.47.

We also examine whether this change in the schedule led voucher holders to move to higher quality buildings. We predict physical structure quality by applying the hedonic coefficients from Section IIIA to data in Dallas on the number of bedrooms, structure type, and structure age (but not building location). In 2010, voucher holders who lived in higher quality neighborhoods had lower structure quality, as would be expected given the existence of a single, metro-wide rent ceiling. We find that for every dollar change in the rent ceiling, structure quality changed by \$0.19, as reported in online Appendix Table 6. This evidence reaffirms that the tilting policy muted the trade-off between unit quality and neighborhood quality. However, this measure does not incorporate the improvements in neighborhood quality that we explore in the next section.

B. *Impacts on Neighborhood Quality*

We assemble data on five measures of neighborhood quality: poverty rate, fourth grade test scores at zoned school, unemployment rate, share of children in families

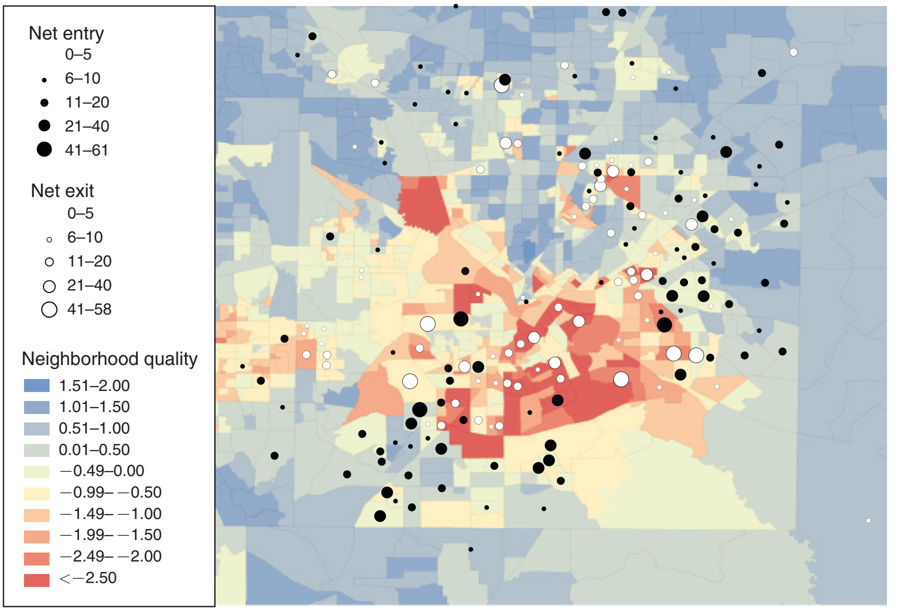

Figure 9. Neighborhood Changes for Dallas Vouchers, 2010–2013

Notes: In 2011, Dallas replaced a single, metro-wide FMR with ZIP code-level FMRs, raising rent ceilings in expensive neighborhoods and lowering rent ceilings in cheap neighborhoods. We construct a neighborhood quality index as an equally weighted sum of tract-level poverty rate, test scores, unemployment rate, share of kids with single mothers, and violent crime rate. The index is normalized to have mean zero and unit standard deviation with respect to the entire Dallas metro area. Black circles indicate increasing number of vouchers in a tract and white circles indicate decreasing number of vouchers in a tract.

with single mothers, and the violent crime rate.[22] We compute a neighborhood quality index, which equally weights all five measures.[23] Figure 9 shows Dallas, with the neighborhood quality index colored from red (lowest) to blue (highest). Voucher holders tend to live in lower quality neighborhoods, usually on the south side of the city. Figure 9 also shows the change in voucher counts at the tract level from 2010 to 2013. A black dot indicates a net increase, a white dot represents a net decrease, and the size of the dot indicates the magnitude of the change. After the policy change, voucher holders exit the lowest quality neighborhoods in the inner city, moving further south and east to better neighborhoods. Figure 9 shows that the improvement in neighborhood quality was broad-based, and not driven by moves to or away from a single neighborhood.

To formally estimate the impact of the change to ZIP code-level FMRs, we use a simple difference-in-differences design with a comparison group of Fort Worth—a nearby city that continued to have a single metro-wide rent ceiling. We construct a

[22] Poverty rate, unemployment, and share of kids in families with single mothers are ACS tract-level data from 2006 to 2010. Test scores are the percent of fourth grade students scoring proficient or higher on state exams in the 2008–2009 academic year at zoned school. Violent crime is number of homicides, non-negligent manslaughter, robberies, and aggravated assaults per capita in 2010, and is calculated over the tract level for tracts in the city of Dallas, and at the jurisdiction level (city or county balance) for suburban voucher residents.

[23] Each component is standardized to have mean zero and unit standard deviation over the Dallas metro area.

balanced panel of voucher holders in the eight affected counties from 2010 to 2013 to mitigate any unrelated composition changes over time.[24] The identifying assumption is that neighborhood quality difference between Dallas voucher tenants and Fort Worth voucher tenants would have been stable absent the policy intervention. We estimate

$$(7) \qquad Y_{it} = \alpha + \delta \mathbf{1}(Dallas_i \times Post_t) + \mathbf{1}(Dallas_i) + \mathbf{1}(Post_t) + \eta_{it},$$

where i indexes households and t indexes years, $\mathbf{1}(Dallas_i)$ is an indicator taking the value one if the voucher holder i is with an affected Dallas housing authorities, and zero if the voucher holder is with a Fort Worth housing authority; and $\mathbf{1}(Post_t)$ is an indicator if the observation is after the policy change became effective. The results are shown in Table 4, where δ shows an intent-to-treat (ITT) improvement of 0.10 standard deviations in neighborhood quality. This estimate is statistically precise, with a t-statistic greater than three using standard errors clustered at the tract level. Of course, neighborhood quality could only improve for tenants who moved. From 2010 to 2013, 46 percent of continuing voucher holders moved units, so the impact estimate for treatment-on-the-treated (TOT) is 0.23 standard deviations.[25]

Table 4 also provides impacts separately for each of the five neighborhood quality measures. We find small and statistically insignificant improvements of 0.09 standard deviation in test scores at zoned schools and 0.05 standard deviation in the neighborhood rate of children living with single mothers. We find medium-sized improvements of 0.19 standard deviation in the neighborhood poverty rate and 0.21 in the neighborhood unemployment rate. The largest improvements are in the violent crime rate, which improves by 0.33 standard deviation. If these relative improvements reflect voucher holders' valuations, then it seems that voucher holders prioritize getting away from high crime areas. This is consistent with evidence from the Moving to Opportunity (MTO) experiment, where treatment households chose tracts with much lower crime rates, less graffiti, and better police response when a call was made (Kling, Ludwig, and Katz 2005).

The timing and distribution of neighborhood choices is consistent with attributing the results in Table 4 to the impact of the policy. Figure 10 shows that neighborhood quality moves in tandem for Dallas and Fort Worth through 2010; beginning in 2011, there is an immediate and sustained increase in Dallas that does not appear in Fort Worth. Online Appendix Figure B.3 shows the distribution of neighborhood

[24] We use a balanced panel to isolate the effects of the intervention on neighborhood quality. During this period, some housing authorities changed the allocation rules for new vouchers. For example, beginning in 2009 the Dallas Housing Authority allocated many of its new vouchers to homeless individuals. These individuals needed other non-housing services and are a very different population from standard voucher holders. Nevertheless, when we analyze impacts for new voucher recipients, they also show improved neighborhood quality after the policy change.

[25] The court settlement that precipitated the policy change also funded voluntary mobility counseling, provided by the Inclusive Communities Project, the organization that filed the lawsuit. There were 303 voucher households who already had conventional vouchers in 2010 and took advantage of these counseling services by the end of 2012. Online Appendix Table 7 shows that households that received counseling showed dramatic improvements in neighborhood quality of 1.17 standard deviations. These large impacts may reflect self-selection or the causal impact of the intervention. If the quality improvement for these 303 households is entirely attributable to the causal impact of mobility counseling (and not to the ZIP code-level FMRs), then our estimates for the impact of ZIP code-level FMRs shrink by about 20 percent.

Table 4—Effect of Tilting Rent Ceilings to ZIP-Level on Neighborhood Quality (*research design: Dallas*)

	Fort Worth (control)		Dallas (treatment)		Differences		Diff-in-Diff (ITT)[e]	Diff-in-Diff (TOT)[f]	Standardized effect[g]
	Pre (1)	Post (2)	Pre (3)	Post (4)	(2) − (1) (5)	(4) − (3) (6)	(6) − (5) (7)	(8)	(8)/SD (9)
Poverty rate[a]	0.174	0.172	0.210	0.199	−0.001	−0.011	−0.009 (0.003)	−0.0210	0.188
Test scores[b]	−0.719	−0.707	−0.494	−0.445	0.012	0.049	0.037 (0.030)	0.0819	0.085
Unemployment	0.096	0.097	0.107	0.104	0.001	−0.003	−0.004 (0.001)	−0.0089	0.208
Single mothers	0.363	0.356	0.381	0.370	−0.008	−0.011	−0.003 (0.004)	−0.0076	0.047
Violent crime[c]	0.0067	0.0066	0.0151	0.0138	−0.0001	−0.0013	−0.0012 (0.000)	−0.0026	0.327
Neighborhood index[d]	−0.700	−0.684	−1.105	−0.986	0.017	0.118	0.102 (0.028)	0.225	0.225
Rent (2010 $)	709	700	796	777	−8	−19	−10 (4.066)	−23	
Observations	7,203	7,038	19,315	19,399					
n moved		3,041		8,899					

Notes: This table shows the neighborhood quality impact of moving from a single, metro-wide FMR in Dallas to ZIP-level FMRs. See Section VIB for details.

[a] Poverty rate, unemployment, and share of kids in families with single mothers are ACS tract-level data from 2006 to 2010.

[b] Percent of fourth grade students scoring proficient or higher on state exams in the 2008–2009 academic year at zoned school. Proficiency rates are standardized to have mean zero and unit standard deviation over block-groups in the Dallas metro area.

[c] Violent crime is number of homicides, nonnegligent manslaughter, robberies, and aggravated assaults per capita in 2010, and is calculated over the tract level for tracts in the city of Dallas, and at the jurisdiction level (city or county balance) for suburban voucher residents.

[d] Index is an equally weighted sum of the five measures, standardized to have mean zero and unit standard deviation.

[e] Intent-to-Treat Estimates. Standard errors for Diff-in-Diff estimate in column 7 are clustered at the tract level and are in parentheses.

[f] Treatment-on-Treated Estimates. Column 7 divided by the fraction of Dallas tenants who moved to a new unit.

[g] Standardized effect is Diff-in-Diff estimate with each measure reoriented so that positive indicates an improvement, divided by standard deviation for all census tracts in the Dallas metro area.

qualities chosen by movers; movers after the policy change appear to have a broad-based monotonic shift away from lower quality neighborhoods and to higher quality neighborhoods. No such change is evident for the control group in Fort Worth.

Averaging across the entire Dallas metro area, average voucher rents are essentially unchanged after tilting the rent ceiling, as shown in Table 4. Given that average neighborhood quality rose, it is somewhat surprising that this policy was budget neutral. The reason for this is that there is heterogeneity in where voucher holders live and they usually live in low-quality neighborhoods. Because they are concentrated in low-quality, inexpensive neighborhoods, the policy would have saved money absent any behavioral response in terms of improved neighborhood quality. Coincidentally, the additional expenditure on improved neighborhoods almost exactly offsets the cost savings from the policy.

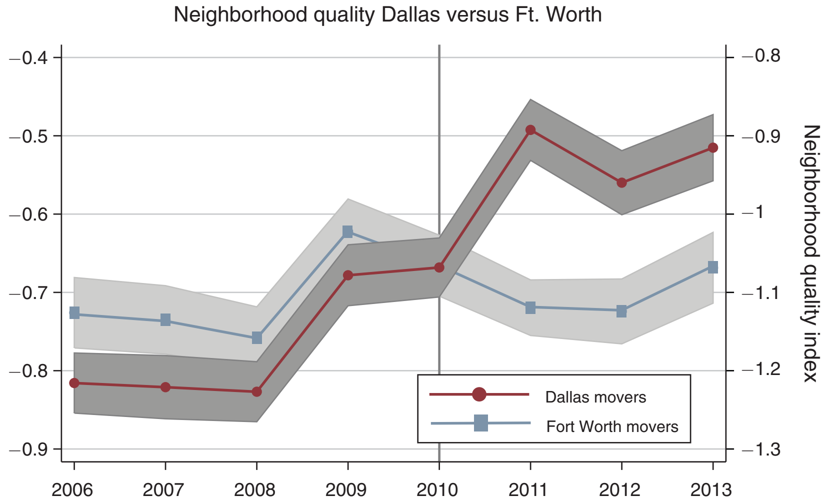

Figure 10. Impact of Tilting: Neighborhood Quality (Time-series)

Notes: In 2011, Dallas replaced a single, metro-wide FMR with ZIP code-level FMRs, raising rent ceilings in expensive neighborhoods and lowering rent ceilings in cheap neighborhoods. We construct a neighborhood quality index as an equally weighted sum of tract-level poverty rate, test scores, unemployment rate, share of kids with single mothers, and violent crime rate. The index is normalized to have mean zero and unit standard deviation with respect to the entire Dallas metro area. The above figure plots the average neighborhood quality for movers in each year in the Dallas metro area and the Fort Worth metro area. The left vertical axis is the quality level of Fort Worth movers, the right vertical axis reports the quality level of Dallas movers, and both axes share the same scale.

C. *Comparing Policies to Improve Neighborhood Quality*

The impact on neighborhood poverty rates for voucher holders of the Dallas policy is substantial in comparison with the uniform increases studied in Section III. We consider three scenarios: (i) a 10 percent increase in the rent ceiling, multiplied by the coefficient from the rebenchmarking estimate, (ii) a shift of FMRs from the fortieth to the fiftieth percentile, and (iii) the Dallas policy. The rebenchmarking yields a precise zero, the shift to the fiftieth percentile yields an imprecise zero, and the Dallas policy yields an improvement which is statistically large and economically significant.[26]

We also compare the neighborhood quality impacts in Dallas to other randomized voucher interventions in Table 5. Voucher holders' access to areas with good schools, low poverty, and low crime has been a major focus of research in recent years (Lens, Ellen, and O'Regan 2011; Horn, Ellen, and Schwartz 2014). Two prominent studies with random assignment of vouchers where the tract-level poverty rate and violent crime rate are available as outcome measures are the Moving to Opportunity (MTO) experiment and voucher random assignment in Chicago (Jacob and Ludwig 2012; Jacob, Ludwig, and Miller 2013). These studies are informative about two types of policy interventions: giving a voucher to someone in public housing and giving a voucher to someone receiving no housing assistance. From largest to smallest, the

[26] The results are shown in a bar graph in online Appendix Figure B.4.

Table 5—Comparison of Policies to Improve Neighborhood Quality

Neighborhood measure	Poverty rate Control	Poverty rate Treat	Violent crime Control	Violent crime Treat	Annual cost (2010 $)	Predicted impact on child income rank
Voucher with ZIP-Level FMR versus Metro-wide FMR						
Tilting rent ceiling (Dallas)	21.0%	18.9%	151	125	−$23	3.1
Voucher versus public housing						
Moving to opportunity experimental	42%	18%	234	128	$2,144	16.8
Moving to opportunity Section 8	42%	28%	234	211	$2,144	6.0
Lottery from Chicago public housing	48%	22%	219	201	$2,144	8.6
Voucher versus no voucher						
Lottery from Chicago private housing	25.7%	24.6%	167	166	$5,299	0.4

Notes: "Treat" is constructed as control mean plus impact estimate for Treatment-on-Treated. Poverty rate and violent crimes per 10,000 residents are tract-level data. *Cost:* Annual cost of Dallas program is from Table 5. Annual cost of a voucher subsidy is equal to 12 times contract rent plus utility allowance minus tenant contribution from Table 1. Annual cost of moving someone from public housing to a voucher is cost of voucher subsidy from Table 1 minus annual ongoing maintenance cost of a public housing unit (estimated as $3,155/year by Finkel et al. 2010). *Predicted Impact on Child Income Rank:* Chetty and Hendren (2017) provides estimates of the cross-sectional relationship between the estimated causal effect of childhood exposure to a county on adult earnings and the poverty rate and violent crimes of a county. We estimate the impact of the poverty rate and the violent crime rate on the income rank of a child whose parents are at the twenty-fifth percentile of the income distribution using their published data. Under the assumption that the cross-county within Commuting Zone coefficients are accurate for the causal impacts of tract-level variation in neighborhood quality, we can calculate the impact of each mobility policy on income of a child who experiences each policy at age 0 and stays in that location until age 20.

Source: Moving to Opportunity results from Table 2, Kling, Ludwig, and Katz (2005). Lottery from Chicago Public Housing from Table 2, Jacob, Ludwig, and Miller (2013). Lottery from Chicago Private Housing from Table V, Jacob and Ludwig (2012).

improvements are largest for the MTO experimental group, who were *required* to move to low-poverty tracts, medium-sized for people leaving public housing with unrestricted vouchers, and zero for unassisted tenants given unrestricted vouchers. The improvements for people leaving public housing are unusually large in part because holders were leaving distressed public housing with a high concentration of poverty.

For each intervention, we construct a cost estimate and summary measure of the change in opportunity for a child affected by the policy. We construct our summary measures as an estimated effect on children's income rank as adult at age 30. Chetty and Hendren (2017) document heterogeneity in intergenerational mobility across US commuting zones. We use estimates from Chetty and Hendren (2017) of the cross-sectional relationship between the causal estimates of a childhood spent in a county and that county's violent crime and poverty rates to generate these predictions. To be precise, we take estimates from Table XII, which reports the results of univariate regressions of the estimated "place effect" on a county characteristic (within commuting zone): for violent crime β^{sd}_{crime} (-1.99) and poverty β^{sd}_{pov} (-1.44). We then calculate the effects of each voucher intervention as

$$\Delta Rank = \frac{\Delta Crime}{\sigma_{crime}} \beta^{sd}_{crime} + \frac{\Delta Pov}{\sigma_{Pov}} \beta^{sd}_{Pov},$$

where $\Delta Crime$ is the treatment and control difference from the intervention in the violent crime rate and σ_{crime} is the standard deviation of the violent crime rate in the

Chetty and Hendren (2017) sample. Similarly, ΔPov is the intervention's effect on tract poverty rates and σ_{Pov} is the standard deviation of poverty.

Our estimates of the causal impact of voucher interventions on children's outcomes make the following assumptions: (i) the child lived in the new location from birth to age 20; (ii) the cross-sectional relationship between the county characteristics and estimates of the causal effect of places from Chetty and Hendren (2017) are accurate for the causal impacts of tract-level variation in neighborhood quality; and (iii) the interventions only affect a child's adult earnings through impacts on neighborhood poverty and violent crime rates. The Chetty and Hendren (2017) results, combined with our assumptions, suggest that tilting the rent ceiling in Dallas with ZIP-level rent ceilings would raise a child's income rank at age 30 by 3.1 percentile points, from the thirty-ninth percentile to the forty-second percentile. This improvement for Dallas is smaller than the predicted improvement for the MTO Experimental group (17 percentage points), about one-half of the impact of offering unrestricted vouchers to public housing residents in MTO, and larger than offering vouchers to unassisted tenants.[27] We approximate the cost of receiving a voucher from public housing with the difference between the average annual cost of a voucher in our sample and an accounting estimate of the per unit cost to maintain the existing public housing inventory (Finkel et al. 2010).[28] Based on these simple cost comparisons, tilting the rent ceiling in Dallas was a cost-effective way to improve opportunity in Dallas.

V. Conclusion

We examine who benefits from two policies designed to improve the neighborhood quality of voucher holders: raising the rent ceiling uniformly and tilting the rent ceiling so that it is higher in high-quality neighborhoods and lower in low-quality neighborhoods. Across two separate research designs we find that increasing the rent ceiling uniformly by \$1 raises voucher rents by roughly \$0.46 with no commensurate improvements in housing or neighborhood quality. In contrast, tilting the rent ceiling in Dallas causes voucher families to move to notably safer and less impoverished neighborhoods at zero net cost to the government. Although tilting the rent ceiling is highly cost-effective and voucher holders move to *better* neighborhoods, the destination neighborhoods are still of a relatively low quality relative to the distribution for Dallas as a whole. Future research should seek to identify other barriers or preferences which affect the neighborhood quality of voucher holders.

[27] This 3 percentage point prediction is if the policy moved children at birth and they stayed in the same neighborhood until age 20. In fact, the improvement neighborhood quality for the MTO experimental group decayed by about 80 percent, so the quality impact of MTO was smaller than the impact of the hypothetical policy considered here which permanently implemented voucher restrictions.

[28] This cost comparison makes no attempt to adjust for housing quality. Also, a more comprehensive cost comparison would take into account the opportunity cost of public housing land and structures, which are not reflected here. The per-family cost of providing a voucher is typically less costly than providing a new public housing unit. For a comprehensive review of studies on the cost of providing voucher and project-based subsidies see Olsen (2008).

REFERENCES

Carlson, Deven, Robert Haveman, Thomas Kaplan, and Barbara Wolfe. 2012. "Long-term effects of public low-income housing vouchers on neighborhood quality and household composition." *Journal of Housing Economics* 21 (2): 101–20.

Chetty, Raj, and Nathaniel Hendren. 2016. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." National Bureau of Economic Research (NBER) Working Paper 23001.

Chetty, Raj, and Nathaniel Hendren. 2017. "The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates." National Bureau of Economic Research (NBER) Working Paper 23002.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106 (4): 855–902.

Chyn, Eric. 2016. "Moved to Opportunity: The Long-Run Effect of Public Housing Demolition on Labor Market Outcomes of Children." https://appam.confex.com/appam/2016/webprogram/Paper16616.html.

Collinson, Robert, and Peter Ganong. 2018. "How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality?: Dataset." *American Economic Journal: Economic Policy.* https://doi.org/10.1257/pol.20150176.

Desmond, Matthew, and Kristin L. Perkins. 2016. "Are Landlords Overcharging Housing Voucher Holders?" *City and Community* 15 (2): 137–62.

Eriksen, Michael, and Amanda Ross. 2013. "The Impact of Housing Vouchers on Mobility and Neighborhood Attributes." *Real Estate Economics* 41 (2): 255–77.

Finkel, Meryl, and Larry Buron. 2001. *Study on Section 8 Voucher Success Rates.* Abt Associates Inc. Cambridge, MA, August.

Finkel, Meryl, Ken Lam, Christopher Blaine, R. J. de la Cruz, Donna DeMarco, Melissa Vandawalker, and Michelle Woodford. 2010. *Capital Needs in the Public Housing Program*. Abt Associates Inc. Cambridge, MA, November.

Glaeser, Edward L., and Joseph Gyourko. 2006. "Housing Dynamics." National Bureau of Economic Research (NBER) Working Paper 12787.

Gordon, Nora. 2004. "Do federal grants boost school spending? Evidence from Title I." *Journal of Public Economics* 88 (9–10): 1771–92.

Greater New Orleans Fair Housing Action Center. 2009. *Housing Choice in Crisis: An Audit Report on Discrimination against Housing Choice Voucher Holders in the Greater New Orleans Rental Housing Market*. New Orleans: Greater New Orleans Fair Housing Action Center.

Horn, Keren Mertens, Ingrid Gould Ellen, and Amy Ellen Schwartz. 2014. "Do Housing Choice Voucher holders live near good schools?" *Journal of Housing Economics* 23: 28–40.

Integrated Public Use Microdata Series, Current Population Survey: Version 5.0. 2010. Minnesota Population Center, University of Minnesota. https://cps.ipums.org/cps/citation.shtml (accessed October 18, 2012).

Jacob, Brian A., and Jens Ludwig. 2012. "The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery." *American Economic Review* 102 (1): 272–304.

Jacob, Brian A., Jens Ludwig, and Douglas L. Miller. 2013. "The effects of housing and neighborhood conditions on child mortality." *Journal of Health Economics* 32 (1): 195–206.

Kling, Jeffrey R., Jens Ludwig, and Lawrence F. Katz. 2005. "Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment." *Quarterly Journal of Economics* 120 (1): 87–130.

Lens, Michael, Ingrid Gould Ellen, and Katherine O'Regan. 2011. *Neighborhood Crime Exposure Among Housing Choice Voucher Households*. New York University Wagner School and Furman Center. New York, February.

National Historical Geographic Information System: Version 2.0. 2011. Minnesota Population Center, University of Minnesota. https://www.nhgis.org (accessed March 23, 2016).

Olsen, Edgar O. 2003. "Housing Programs for Low-Income Households." In *Means-Tested Transfer Programs in the United States*, edited by Robert A. Moffitt, 365–441. Chicago: University of Chicago Press.

Olsen, Edgar O. 2008. "Getting More from Low-Income Housing Assistance." Hamilton Project Discussion Paper 2008–13.

Patterson, Rhiannon, Michelle Wood, Ken Lam, Satyendra Patrabansh, and Gregory Mills. 2004. *Evaluation of the Welfare to Work Voucher Program: Report to Congress*. Abt Associates Inc. Cambridge, MA, March.

Quadel Consulting Corporation. 2001. *Voucher Program Guidebook: Housing Choice.* US Department of Housing and Urban Development, Office of Public and Indian Housing. Washington, DC, 2001.

Susin, Scott. 2002. "Rent vouchers and the price of low-income housing." *Journal of Public Economics* 83 (1): 109–52.

US Department of Housing and Urban Development. 2000. "Rule to Deconcentrate Poverty and Promote Integration in Public Housing." *Federal Register* 65 (247): 8123–29.

US Department of Housing and Urban Development. 2014. "Public and Indian Housing Tenant-Based Rental Assistance: 2015 Summary Statement and Initiatives (Dollars in Thousands)." https://www.hud.gov/sites/documents/FY15CJ_TNT_BASED_RNTL_ASST.PDF.