Supplementary Information for

Longer Trips to Court Cause Evictions

David A. Hoffman, Anton Strezhnev

David A. Hoffman.
E-mail: dhoffman@law.upenn.edu

This PDF file includes:

Supplementary text
Figs. S1 to S22
Tables S1 to S5
SI References

Supporting Information Text

1. Description of data pre-processing

Our pre-processing of the data starts with identifying those entries in the docket that we can unambiguously geocode to a particular building in Philadelphia. The complete dataset that we obtained from Philadelphia Legal Assistance (PLA) contains a total of 339,172 docket entries from 2005 to 2021. We first subset down to only those entries where an address of a building can be cleanly extracted and merged to our building ownership dataset from Pew. This requires us to be able to unambiguously parse a building number and street name from the docket address listing, a challenging task since there is no consistent formatting for the defendant's address in the docket. Abbreviations are not standardized, unit numbers sometimes precede and sometimes follow the building number and street name, and occasionally multiple addresses are listed. While PLA was able to address some of these issues via an initial cleaning step applied to the raw docket address, we are unable to merge with the Pew building data without an additional pre-processing step. To obtain a consistent address formatting, we use the `postmastr` R package* to extract the building number, street direction (e.g. North, South), street name and street suffix (e.g. Avenue or Drive) from each docket entry's address field. Each building in the Pew dataset can be uniquely identified by this combination of fields. We remove all entries with addresses for which we could not identify a house number, street name or street suffix. This step leaves 236,300 docket entries for 62,333 unique buildings. Each of these entries could be matched to a building in the Pew dataset.

Our next step was to obtain commuting distance and commuting time estimates using the Google Maps API. We queried the API for each of our 62,333 addresses and obtained an estimated travel time between the building address and the Philadelphia Municipal Courthouse via public transit for a weekday: Wednesday, May 11, 2022. We set the arrival time to 1:00pm. We also obtained an estimate of weekend commuting time for Sunday May 15, 2022 with the same arrival time.† We were unable to obtain a weekday or weekend commuting time for 2,775 docket entries representing 896 buildings. An additional 7 buildings/42 docket entries returned zero results for our weekend commuting time query. We further verified that the address returned by Google Maps contains a building number, street name and street suffix to avoid cases where the geolocation was incomplete. After this pre-processing step, we are left with 232,533 docket entries involving 62,101 unique buildings. 42 of these docket entries have missing weekend commute information but do have weekday commute data.

For each of these buildings, we obtain the latitude and longitude using the Google Maps Geolocation API and use a point-in-polygon spatial join to match each property to its 2010 Census Block and 2010 Census Tract. We match each docket entry to the census block-level data from the 2010 U.S. census on total population, total Black population, total white population and total Hispanic population. We also match by census tract each docket to 2015 American Community Survey data on median income and median contract rent. Because some buildings are geolocated to tracts without income or rent data while others were located in blocks with no recorded population in 2010, we are forced to drop some additional observations. After merging with the census and ACS data, we are left with 229,801 complete docket entries representing 61,852 unique buildings. Using our Pew data, we identified 25,919 unique landlords for these 61,852 buildings.

Finally, we identify each unique defendant in each docket entry and determine the first outcome for that defendant (that is, the first judgment registered in the docket after any continuances). We note that in a number of eviction proceedings with multiple defendants, outcomes can often differ across defendants—such as when only one of the tenants appears in court while the other defaults. As such, our outcomes are measured at the level of the individual defendant rather than the docket as a whole. We focus on only those cases with scheduled hearing dates prior to 2022. For these eviction cases, we are able to extract the names of defendants and identify the outcomes in proceedings for 291,308 unique defendants in 229,489 total docket entries covering 61,799 unique buildings and 25,895 landlords. We further remove cases where the date of the final outcome could not be identified or was incorrectly recorded as prior to the date of the first hearing. We also drop all of those cases that had incorrectly recorded filing dates (such as filing dates after the first hearing date). This left 283,812 unique defendants in 223,840 dockets across 61,014 unique buildings and 25,626 landlords. 210,052 of these cases are prior to the Covid-19 pandemic and do not involve public housing. This is the primary dataset that we use to generate descriptive statistics for case outcomes and default rates over time.

We might be concerned that our pre-processing procedure results in a sample of buildings that is unrepresentative of our wholy excludes segments of the city. While it is not possible to directly compare the included versus excluded cases — since the latter by definition cannot be reliably geolocated – we can compare our final dataset of properties to a benchmark. We use the complete property and landlord dataset provided to us by Pew. Each property in the dataset is geolocated to a census tract using the latitude and longitude coordinates associated with the property's unique parcel identifier in the Philadelphia Office of Property Assessment (OPA) property and assessment history dataset. Of the 149,213 property addresses in the Pew dataset, we were able to geolocate 148,974 using the most recent OPA dataset (released June 7, 2022). Likewise, 60,962 of the 61,014 buildings in our dataset could be located in the OPA data. Figure S1 maps the number of properties in the Pew dataset and in our final dataset by census tract. The third panel plots, by census tract, the share of properties found in the Pew dataset that appear in our final dataset. Overall, while there does appear to be some regional variation in the proportion of properties from the Pew dataset that are captured in our data, this likely reflects both variation in the composition of property types across the city — with more commercial rather than residential properties located in the downtown region — and the prevalence

*Available at `https://slu-opengis.github.io/postmastr/`.

†Due to minor issues with API requests occasionally timing out, 16 of the original queries for May 15th were unsuccessful and we re-queried the API for these addresses on the subsequent Sunday May 22nd. To examine the stability of the reported commuting times from the API across similar days and arrival times, we also re-queried 100 randomly selected addresses from the data. 96% of the reported times were identical from May 15th to May 22nd and the largest discrepancy was only 63 seconds.

of eviction in general. We see no obvious outlier regions that have been entirely omitted as a result of our data processing procedure.

For our analysis of default judgments, we further subset down to cases where the date of the first outcome is the same as the date of the first hearing, removing any continuations or deferrals. This results in a final dataset of 245,946 defendants across 194,293 eviction dockets. Of these, we split the dataset by public and non-public housing and into pre- and post-Covid periods. Our largest and primary dataset, the pre-Covid non-public housing subset, consists of 232,709 defendants across 181,958 eviction dockets encompassing 53,578 unique buildings and 23,730 landlords. We have 6 fewer buildings and 35 fewer defendants when removing those observations with missing weekend commuting time data. Our post-Covid non-public housing dataset consists of 2,895 unique defendants across 2,258 eviction dockets encompassing 1,791 unique buildings and 1,233 unique landlords. The pre-Covid public housing dataset consists of 9,389 defendants across 9,310 eviction cases involving 2,556 unique properties.

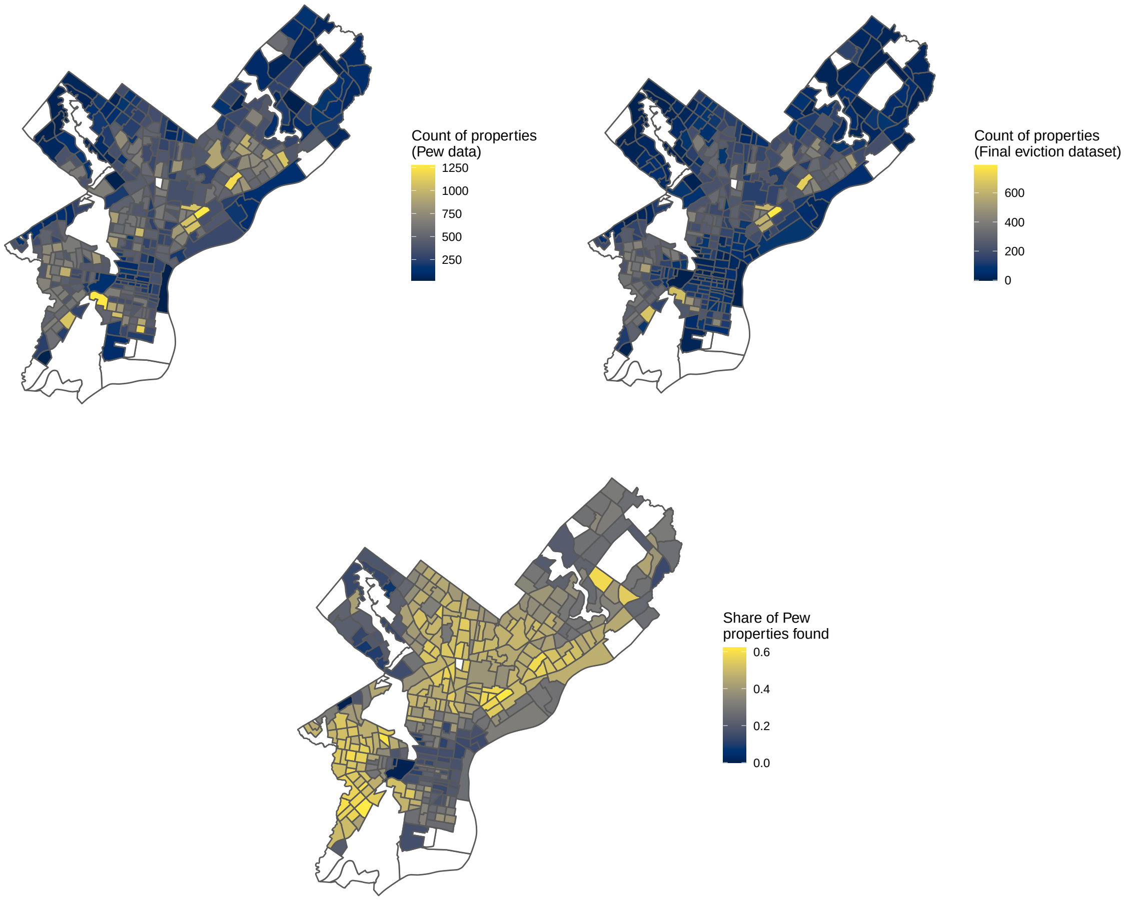

Fig. S1. Distribution of buildings by census tract: Pew dataset compared to our final dataset. 148,974 unique property addresses in the Pew dataset. 60,962 unique properties from the post-processed eviction docket. Census tracts with fewer than 25 properties in the Pew dataset omitted.

2. Sensitivity analysis

We evaluate the sensitivity of our main specification to potential unobserved confounding using the method recommended by Cinelli and Hazlett (1). This approach considers a hypothetical confounder specified by two sensitivity parameters which we vary: the partial R^2 of the relationship between the confounder and commuting time and the partial R^2 of the relationship between the confounder and default.

Figure S2 shows the contour plot of the adjusted point estimates under varying degrees of treatment-confounder and outcome-confounder relationships. The black point at (0,0) denotes our treatment effect estimate under the assumption of conditional ignorability. The red contour line captures all of the possible combinations of treatment-confounder partial R^2 and outcome-confounder partial R^2 that would lead the coefficient to be driven to 0. The "robustness value" is 3.5 percent. In other words, any hypothetical confounder that would eliminate our results would need to explain at least 3.5 percent of the residual variance in both the treatment and the outcome. While this may seem small, it is actually quite a bit larger than the residual variance explained by some of our most salient confounders. Eviction court outcomes are extremely noisy and, as we show in Appendix 4, even highly relevant census-block-level demographic covariates only explain a fraction of the variation. To aid in interpreting the sensitivity contours, we conduct two benchmarking exercises. We first consider a confounder that has an association with treatment and an outcome no larger than that of all of our tract and block-level demographic covariates (income, rent and racial demographics). Such a confounder would explain a reasonable fraction of the residual treatment but, it turns out, very little of the residual outcome. As shown in Figure S2, such a confounder would be insufficient to reduce our point estimate to 0. Likewise, a confounder with a strength comparable to whether a building is an apartment would also not reduce our point estimate to zero. We consider our results fairly robust to potential unobserved confounding as we do not believe a potential confounder would be likely to explain away much of the outcome variance since even our strongest covariates can explain barely 0.2 percent of the residual variance. As a result, any potential unobserved confounder would need to have a very large association with our treatment to induce bias sufficient to eliminate our treatment effect.

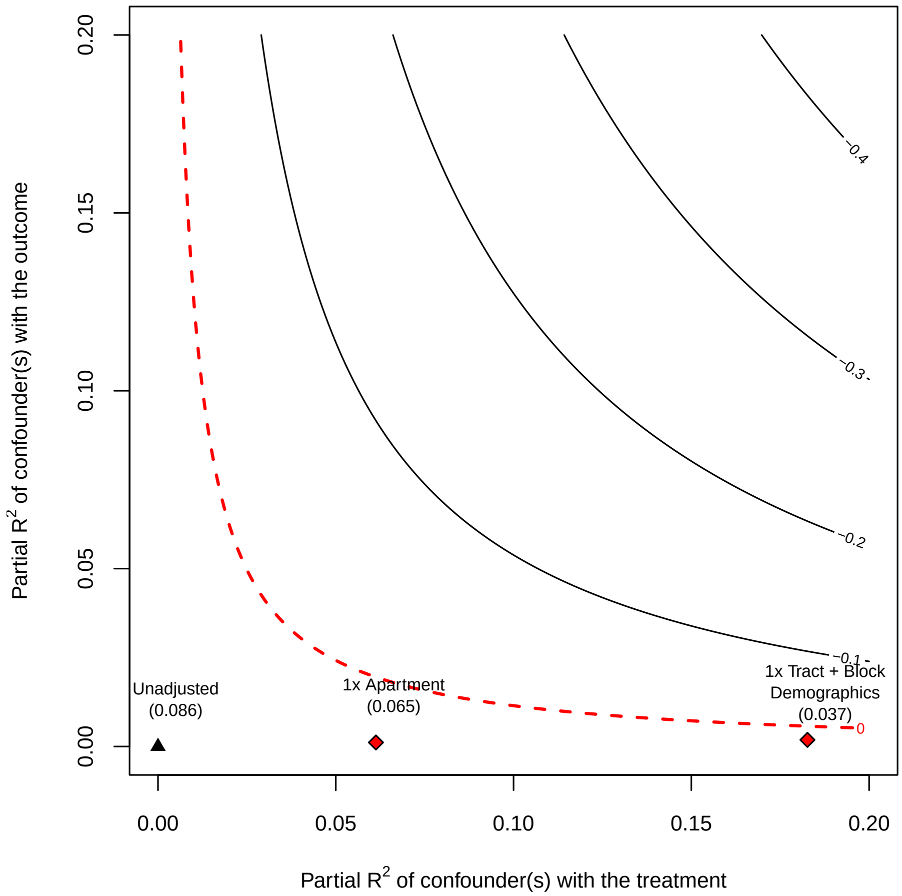

Fig. S2. Sensitivity analysis under hypothetical confounding. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint, and whether the building is classified as an apartment. Month-year fixed effects included.

3. Reopening petitions

How often are default judgments reopened? While not all default judgments are final, we find that defendants rarely file to reopen and that petitions to reopen default judgments are not granted in the majority of cases. As a result, we consider the large majority of defaults that we observe to be final and that there is a direct connection between default judgments and actual evictions.

Previous analyses of petitions to reopen in the Philadelphia Municipal Court provide a window into the reasons behind tenant defaults. Eisenhard et. al. analyzed a sample of 430 cases in which tenants had filed a motion to reopen a default (2). They coded those motions for the reasons offered for the original default and found that the "five most common reasons, which account for 70% of the total, were lack of notice, medical issues, childcare problems, lateness, and difficulty finding the courthouse and courtroom" Notably, and consistent with our findings on the importance of commuting time, a plurality (20%) of defaulting tenants stated that they were merely late—often just a "few minutes"—to court.

We extract from each docket entry in our dataset whether a defaulting tenant filed a notice to reopen and whether that notice to reopen was granted by the court. Among the 232,709 defendants in our pre-Covid, non-public housing dataset (excluding continuances and deferrals), we observe 99,283 unique defaulting tenants. Only 9,511 of these tenants filed a notice to reopen (about 10%). The court granted 3,462 of these notices (about 36%). We find a negligible but slightly positive trend in the share of default judgments that result in petitions to reopen. Figure S3 plots the general trend in the share of reopen petitions and the share of reopening petitions that are successful. We find a small positive trend in the share of default judgments that result in petitions to reopen: the share has risen by about 3 percentage points from 2005 to 2021 with recent reopening petition rates closer to 11 percentage points. The success rate appears to oscillate slightly over time with between 30 and 40 percent of these petitions being granted. We also find no relationship between our primary variable of interest, transit time, and either the probability of filing a reopening petition conditional on default or the probability of a successful reopening petition conditional on filing (Figure S4).

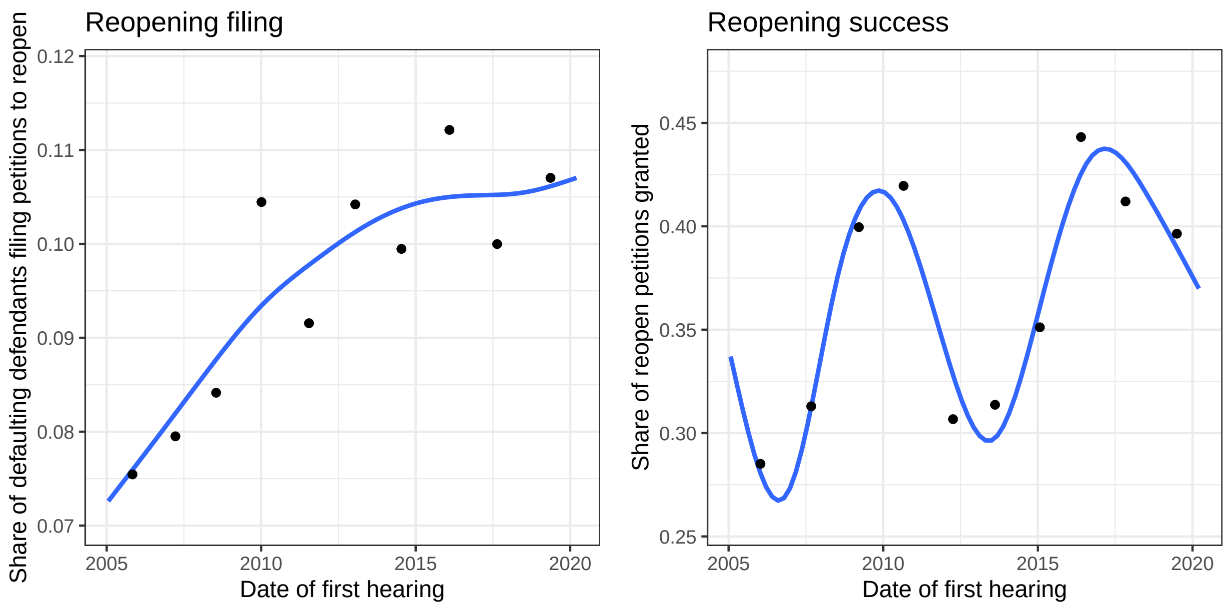

Fig. S3. Rate of reopening petitions and reopening petition success over time. 99,283 pre-Covid non-public housing defendants without continuances/deferrals in their first hearing. 9,511 filed reopening petitions. Lines are cubic-spline smoothers

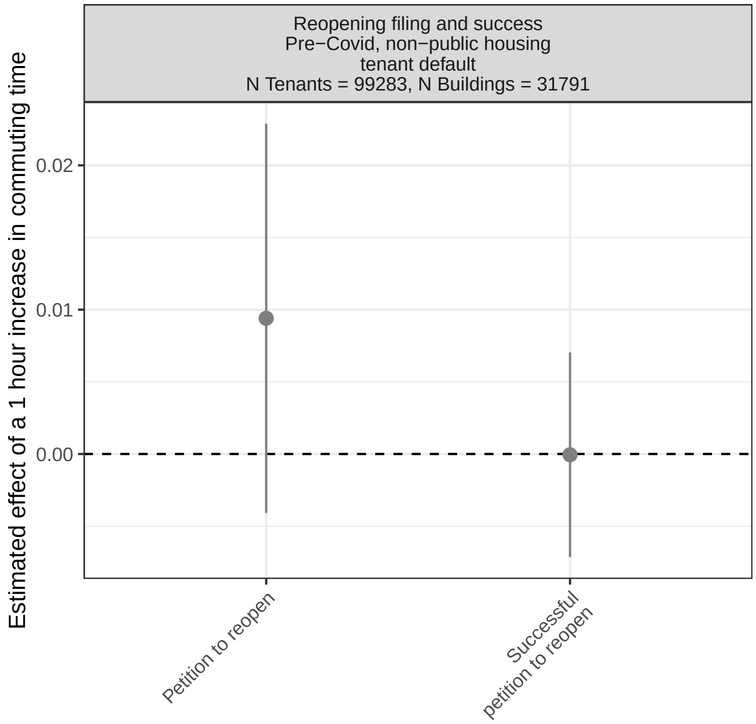

Fig. S4. Estimated average effects of a 1 hour increase in commuting time on probability of filing a reopening petition/filing a successful reopening petition. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint, and whether the building is classified as an apartment. Month-year fixed effects included in all regressions.

4. Additional descriptive statistics

Income. Census tract income levels are correlated with commuting time. Figure S5 plots the relationship between (logged) census tract median income and median commuting time in each census tract. Richer areas of the city tend to be located further away from the city center, though the relationship is curvilinear, with the *most* wealthy tracts having lower average commuting times. We find that default rates also vary by income, though the differences across tract are rather small.

Race and ethnicity. We also observe a relationship between neighborhood race and ethnicity and default rates. On average, census blocks with a larger share of white residents see higher default rates. The relationship is curvilinear, with a steady increase in default rates from about 37.5 percent in census blocks with zero white residents to about 43 percent in blocks that are majority white. Race also is associated with commuting time. Majority white neighborhoods are, on average, located further away from the city center. Conversely, majority black neightobrhoods tend to be closer to the city center. Similar patterns are observed for the percentage of Hispanic residents as well (Figure S8)

Building characteristics. While the Pew data contains a variety of covariates related to the characteristics of the buildings, very few showed a clear relationship with our treatment of interest. One potential confounder that we did notice, however, was whether the building was a multi-unit apartment building (as opposed to a row house or single family dwelling). In previous work, we identified apartment buildings as being more likely to adopt lease templates with terms that are adverse to tenants (3). In particular, almost all templates used by apartments included waivers of notice for eviction proceedings.

We find a bimodal distribution of apartment buildings in the city — basically all buildings in our data that are within 10 minutes of the courthouse (downtown) while essentially none are in the 25 minute to 50 minute zone (primarily north and west Philadelphia where rowhouses are more common). Beyond 50 minutes, however, we see a larger share of apartments, particularly in the northwest of the city. We also find that apartments, on average, have slightly higher default rates – by about 3 percentage points. Further work might explore the determinants of this finding, which could include deficits in the service of process in particular buildings.

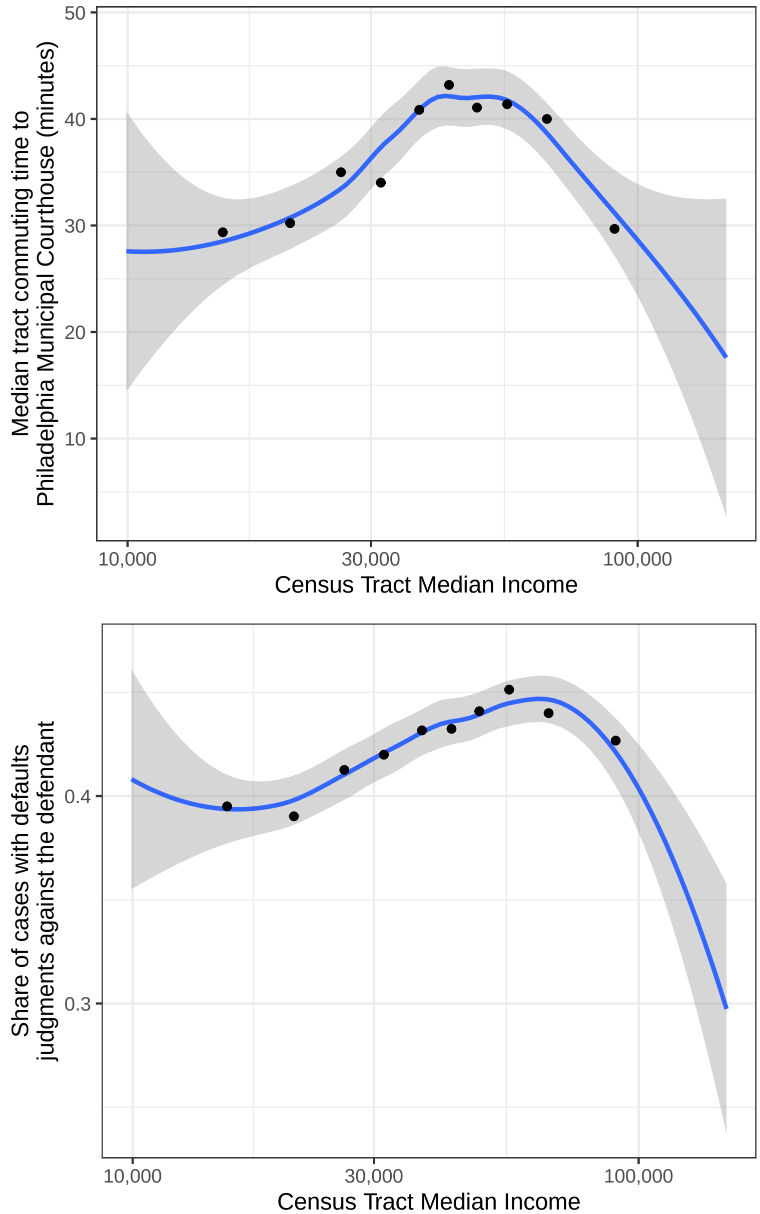

Fig. S5. Relationship between default rates, median commuting times, and census tract median income. N = 232,709 non-public housing tenants (pre-Covid), 365 census tracts. Tracts with fewer than 10 evictions omitted. Plots are loess-smoothed estimates of the regression function.

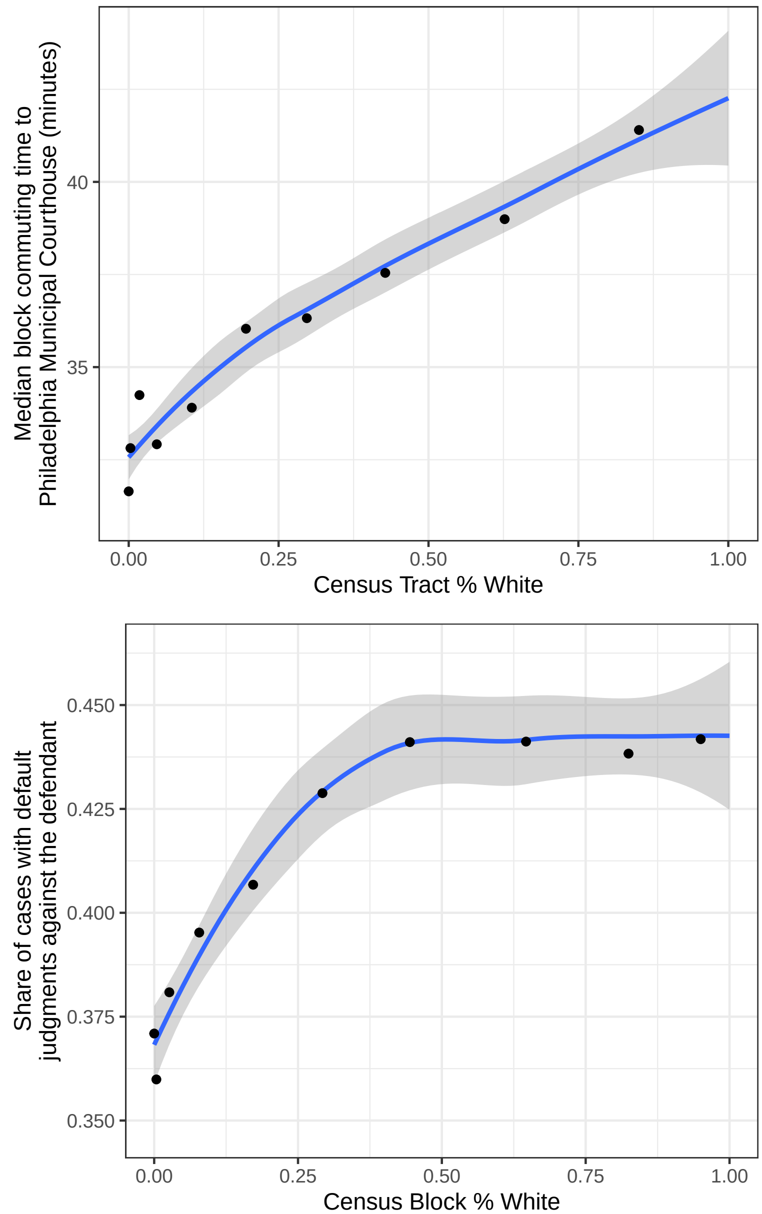

Fig. S6. Relationship between default rates, median commuting times, and census block percent black. N = 232,709 non-public housing individuals (pre-Covid), 5,581 census blocks. Blocks with fewer than 10 evictions omitted. Plots are loess-smoothed estimates of the regression function.

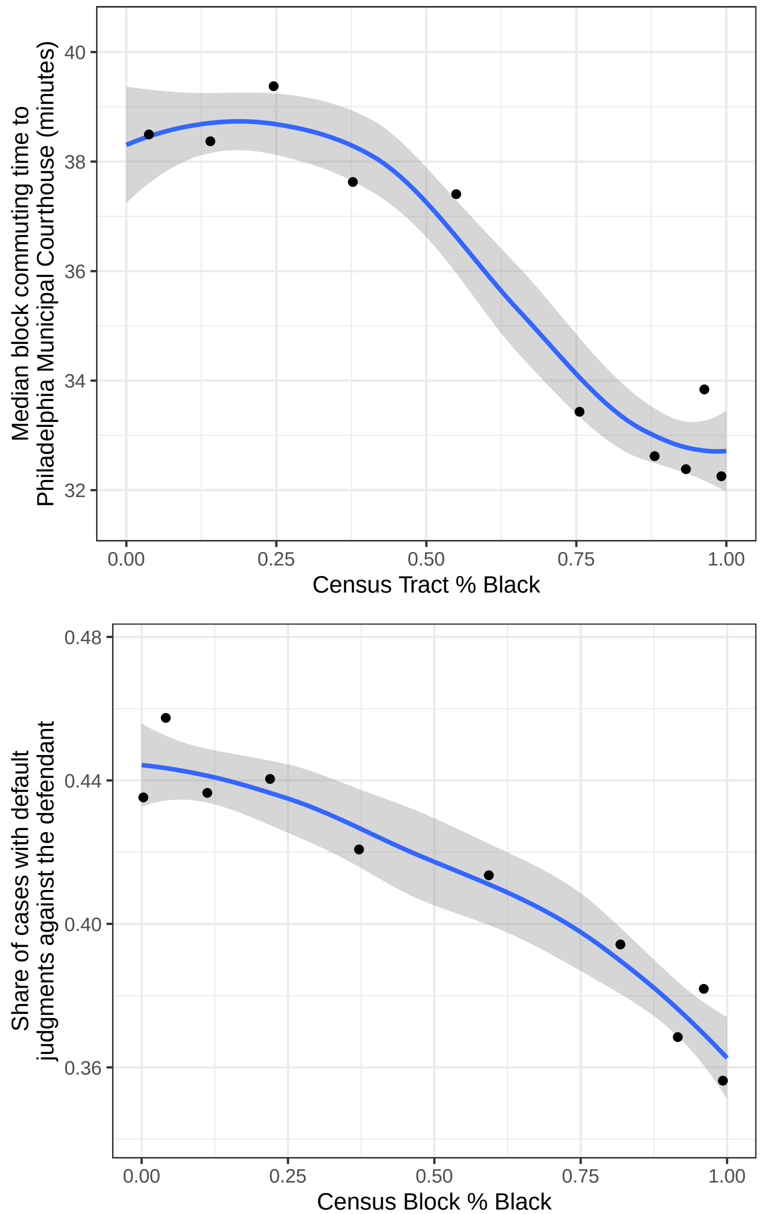

Fig. S7. Relationship between default rates, median commuting times, and census block percent white. N = 232,709 non-public housing individuals (pre-Covid), 5,581 census blocks. Blocks with fewer than 10 evictions omitted. Plots are loess-smoothed estimates of the regression function.

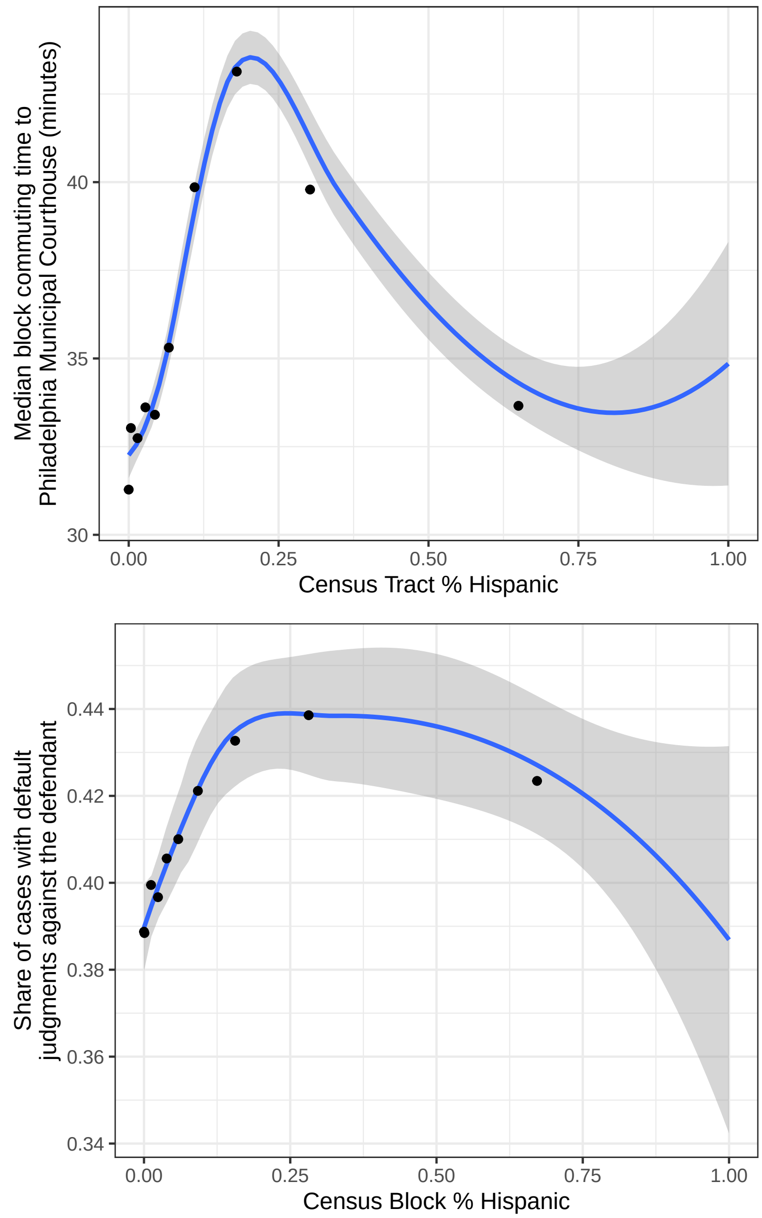

Fig. S8. Relationship between default rates, median commuting times, and census block percent hispanic. N = 232,709 non-public housing individuals (pre-Covid), 5,581 census blocks. Blocks with fewer than 10 evictions omitted. Plots are loess-smoothed estimates of the regression function.

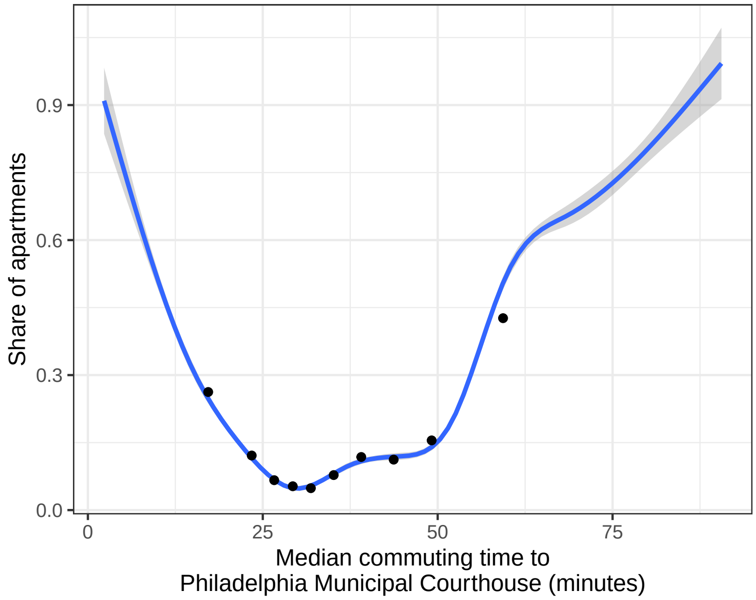

Fig. S9. Relationship between commuting times and prevalence of apartments. N = 232,709 non-public housing individuals (pre-Covid), 53,578 buildings. Plots are cubic spline-smoothed estimates of the regression function.

5. Weekday vs weekend commuting times

We find some notable variation between weekday and weekend commuting times as reported by the Google Maps API. This section provides some additional descriptive statistics on the magnitude of the differences and what regions of the city appear to be driving this variation. Figure S10 plots the histogram of deviations between weekday and weekend commuting times. On average, there is not a meaningful divergence between the two. Indeed, for 31 percent of defendants in our dataset, the difference is exactly zero. However, for a subset of properties, we see differences between weekday and weekend commute ranging from -20 minutes (weekday commutes being *easier*) to 20 minutes (weekday commutes being *harder* than weekend commutes). The second panel of S10 plots the geographic distribution of these deviations. Around the downtown region and slightly around it, we see only minimal variation - at most 5 or 10 minutes in either direction in some areas. The contours of the two subway lines are also visible and slightly negative, likely reflecting the more consistent operating hours of the metro during the weekday versus the weekend. The largest positive and negative gaps appear in the range of properties around 10 to 15 kilometers from the city center in the north of the city. This is consistent with the general geography of Philadelphia transit, where transit quality begins to deteriorate at the edges of the service lines.

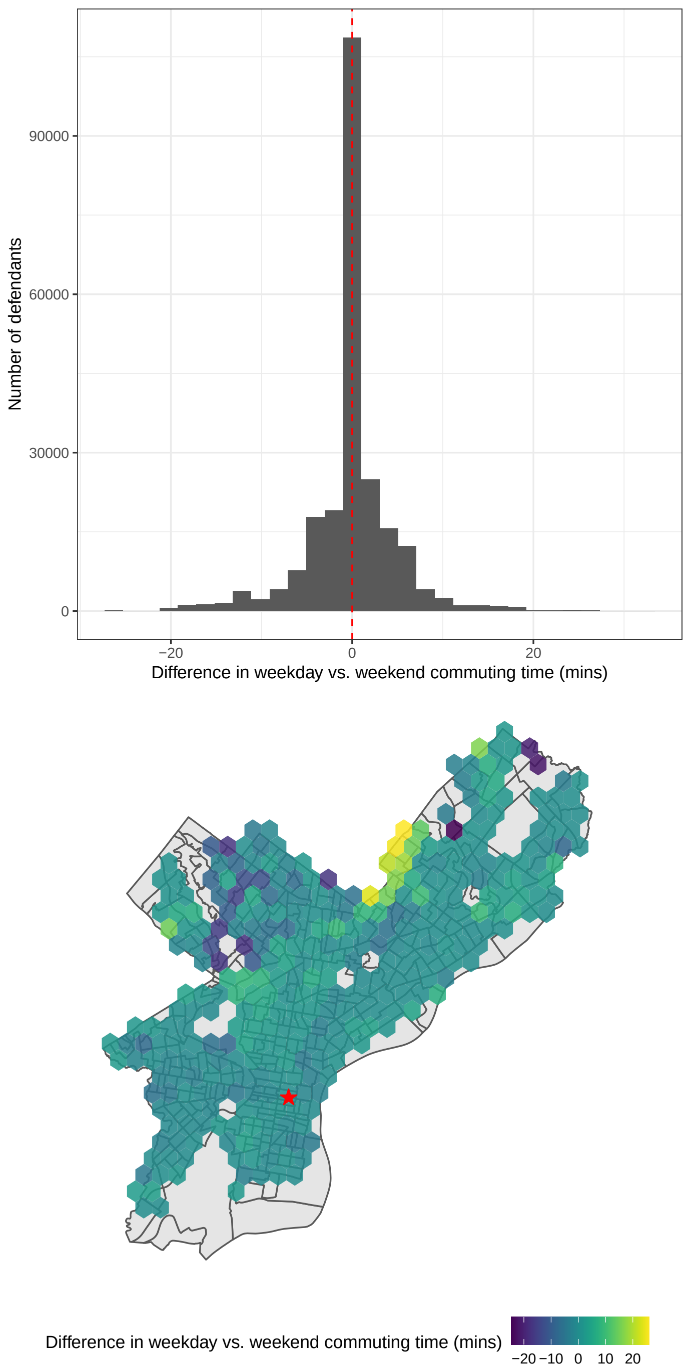

Fig. S10. Distribution of the gap between estimated weekday and weekend commuting time. 53,572 unique buildings, 181,929 eviction proceedings, 232,674 defendants. Transit times estimated using Google Maps Distance Matrix API.

6. Effects of commuting time on judgments by agreement and complaint withdrawals

While commuting time may increase the proportion of default judgments, it is important to also understand what outcomes are being *reduced* as a consequence of distance and the extent to which this is associated with fewer actual evictions of tenants. If it were the case that all defaults due to distance would simply have been judgments against the tenant, then it would be inaccurate to say that distance causes evictions — the tenant would still lose the case by default or in person. To address this, we further estimate the effect of commuting time on the two main alternative outcomes that we observe in Philadelphia evictions: judgments by agreement (JBA) and withdrawals.

To start, we note that after entry of judgment (however derived) the landlord must wait 10 days in the eviction court before being permitted to file for a writ of possession, and then an additional 11 days to file for a "alias writ of possession", which enables them to schedule a lockout through the court staff. When the lockout happens, it is noticed on the docket as "alias writ served" meaning that the tenant was either forcibly removed from the premises or prevented from entering. Advocates tell us that for JBAs and defaults alike, landlord attorneys will typically file a writ of possession and alias writ of possession immediately (assuming the JBA provides for it). This enables the landlord to save time later.

Although only about a quarter of the cases in our data result in the filing of a served alias writ of possession or a lockout eviction, it is useful to characterize the likelihood of this extreme outcome across the different judgments in order to understand the relative benefits of avoiding default compared to a JBA or a case withdrawal. Figure S11 plots the proportion of judgments that resulted in a served alias writ across each of the case outcomes in our data. For the three primary outcomes – withdrawal, JBA and default – we see a clear ordering of preferences. Withdrawals constitute about 20 percent of all outcomes and are clear victories for the tenant – the landlord rarely refiles or pursues the case any further and we see only about 1 percent of all withdrawals resulting in a subsequent served alias writ. JBAs result from private agreements between landlords and tenants at the courthouse. Judgments by agreement are the second most common outcome (34% of cases) and while they may sometimes codify a judgment of possession in favor of the landlord, they often will stop short of such an outcome. (4) analyzes an analogous process in Massachusetts and finds that some two-thirds of agreements result in an agreement where the tenant could stay if certain conditions were met. While not all of these negotiated judgments by agreement result in such a settlement for the tenant, they are nevertheless an option that in most cases preferable to default. However, a recent study by the Reinvestment Fund, which manually coded 2,000 JBAs from Philadelphia in 2018-2019, found a more complicated picture (5). In that sample, approximately 30% of JBAs resulted in an alias writ served (i.e., a lockout) within a year of their filing. This is consistent with our findings as well in S11. An additional 10% had a breach noted, without a subsequent lockout, which might imply that the tenant left the property "voluntarily." The remaining approximately 60% either noted that the judgment was satisfied (i.e., the tenant paid their rent and stayed), there was no further legal action, or the case was otherwise withdrawn. The post-judgment outcomes were strongly associated with access to counsel. The study found that while only 19% of all tenants simply could move out without paying anything, this proportion rose to 63% among tenants with counsel. This finding accords with the literature, finding that JBAs generally result in relatively bad outcomes for unrepresented tenants (6).

Unfortunately, we are unable to directly code for the outcomes in JBAs as they are rarely clearly reported in the digitized docket entry in a way that can be parsed via standard text analysis methods. However, we do note that despite the fact that JBAs often result in outcomes only slightly better for tenants than a default, we still observe fewer lockout evictions for JBAs than among default judgments (30 percent versus 35 percent). This is suggestive of some improvement in tenant outcomes under the JBA process though it of course does not guarantee that a tenant receiving a JBA instead of a default will necessarily be better off. Moreover, we find that even among those cases where a lockout eviction did occur, the time from decision to alias writ served was longer for JBAs. Figure S12 plots the mean number of days between the final judgment and the serving of an alias writ of possession among cases that ultimately resulted in lockouts across different types of initial judgments. Comparing defaults to JBAs, we find that the average post-default lockout takes about 61 days while post-JBA lockouts take about 79 days. This two and a half week difference is generally consistent with the idea that JBAs, on average, are able to buy tenants additional bargaining time relative to a taking default, even if the ultimate outcome is an eviction.

Given that we observe some improvements in post-judgment outcomes among JBAs and withdrawals, it is useful for us to characterize the extent to which each of these outcomes substitutes for defaults. Figure S14 replicates our analysis from the main text but looking instead at the probability of receiving a judgment by agreement or a withdrawal. The estimates mirror the findings for default – longer commute times reduce the probability that a tenant will receive a JBA or that the case will be withdrawn by the landlord. This result holds across all of our adjustment strategies and the placebo results are likewise consistent with the absence of an effect on withdrawals/JBAs during periods where we expect no effect on defaults.

However, it is difficult for us to disentangle whether withdrawal or JBA is the more likely counterfactual outcome for these distance-driven defaults. When we estimate the effect of distance on JBAs alone, we find that the negative effect sharply attenuates to statistical insignificance after adjusting for our confounders (Figure S15). Conversely, the negative effect of distance on withdrawals only appears *after* adjusting for confounders (Figure S16). We think it is possible that confounding is exaggerating the negative effect of distance on JBAs while masking the negative effect on withdrawals. Conditioning on weekend commute times also eliminates the effect on judgments by agreement while the effects on withdrawals remain in three of the four designs. Both sets of estimates fall short of statistical significance after adjusting for landlord fixed effects. Overall, we would conclude that defaults due to distance do not *perfectly* substitute for withdrawals. The mixed results from figures S15 and S16 strongly suggest that effects on these alternate outcomes are likely heterogeneous and that this heterogeneity may be driven by variation in how willing a landlord is to pursue a particular case. Unobserved differences among landlords – which are adjusted for in our landlord fixed-effects regressions – may explain the particular choice of litigation strategy. Landlords that file weaker cases may prefer to withdraw any complaints where the tenant is able to be present while those that file stronger cases may choose to negotiate a JBA.

Lastly, we provide a further note regarding the difficulties in analyzing post-judgment outcomes due to the presence of unobserved negotiations between landlord and tenant and the bias that can be induced by conditioning on a collider (7). Because both parties generally want to avoid a lockout eviction, tenants will often move out voluntarily and comply with judgments against them rather than force the landlord to pursue further post-judgment remedies. Nevertheless, there exist a handful of cases where a defaulting tenant is either unwilling or unable to vacate the premises and the landlord serves an alias writ and forces the tenant to leave. This tenant is likely very different on a variety of unobservable characteristics from the marginal tenant that gets pushed into default due to commute but who would otherwise appear in court.

We can think of our sample as being comprised of two types of tenants – those who would default irrespective of commuting time and those who would default due to having a longer commuting time but who would appear in court if that time were shorter. While the former appear in the set of default judgments across all commuting times, the latter only appear as default when commuting times are long. Because of this variation in composition, the association between commuting time and *alias writs served* may, paradoxically, be zero or negative *even if* shorter commuting times reduce default judgments because the types of individuals who are moved to take a default due to distance are also *less likely* to be the types of tenants who refuse to move out after a judgment. In other words, tenants who would receive a JBA or case withdrawal instead of a default due to a shorter commute are the type that also would not get a lockout eviction under any outcome. We find evidence consistent with this "collider bias" phenomenon. Figure S13 shows that although the gap between the rate at which alias writs are served across our three main judgments remains relatively constant across commuting time, there is an interesting slight *negative* association between commuting time and the propensity of an alias writ served among the default judgments. This is consistent with a pattern where defaults under short commutes are mostly the "hard" cases while those under long commutes are a mixture of "hard" cases and those more marginal tenants for whom commuting time is salient enough to cause them to default.

We further show in Figure S17 that when we replicate our analyses from the main text looking at the probability of an alias writ served, we find no consistent evidence of an average effect of commuting time on the probability of an alias writ served. Moreover, we find strong positive "effects" in our public housing placebos, suggesting that there may be some additional unobserved confounding in the post-judgment outcomes and commuting distances, likely driven by post-judgment interactions between landlord and tenant. These results are complemented by the analysis in Figure S18 which generally finds a negative association between commuting time and the probability of an alias writ being served *within* the subset of cases that receive a default judgment.

As the main text describes, the City of Philadelphia, starting in 2018, rolled out a series of changes to the in-court representation process intended to level the playing field between tenants and landlords, including "navigators" and free legal services. We would anticipate that these changes will result in better outcomes for those tenants who show up to court going forward.

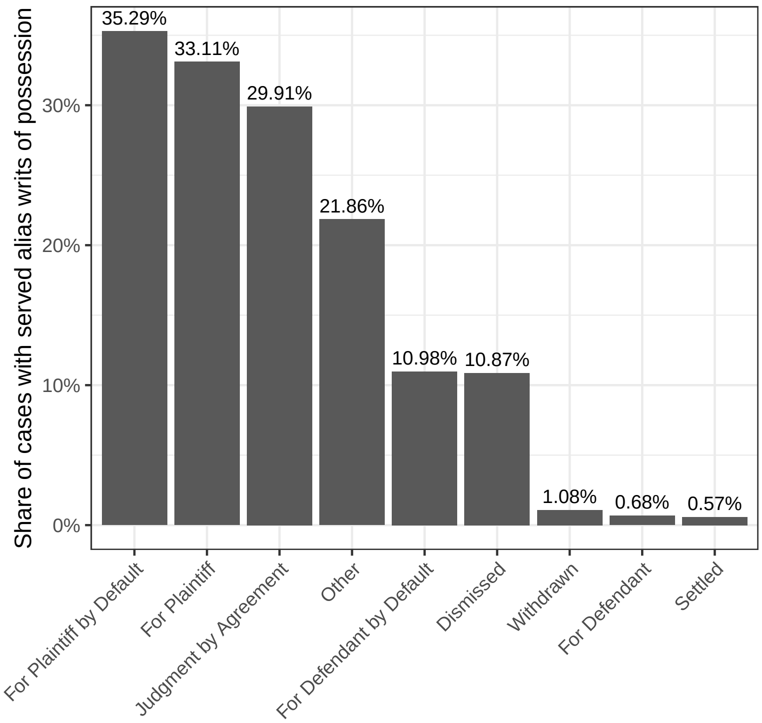

Fig. S11. Proportion of cases with alias writs served: 2005-2021. 223,840 eviction proceedings, 283,812 defendants

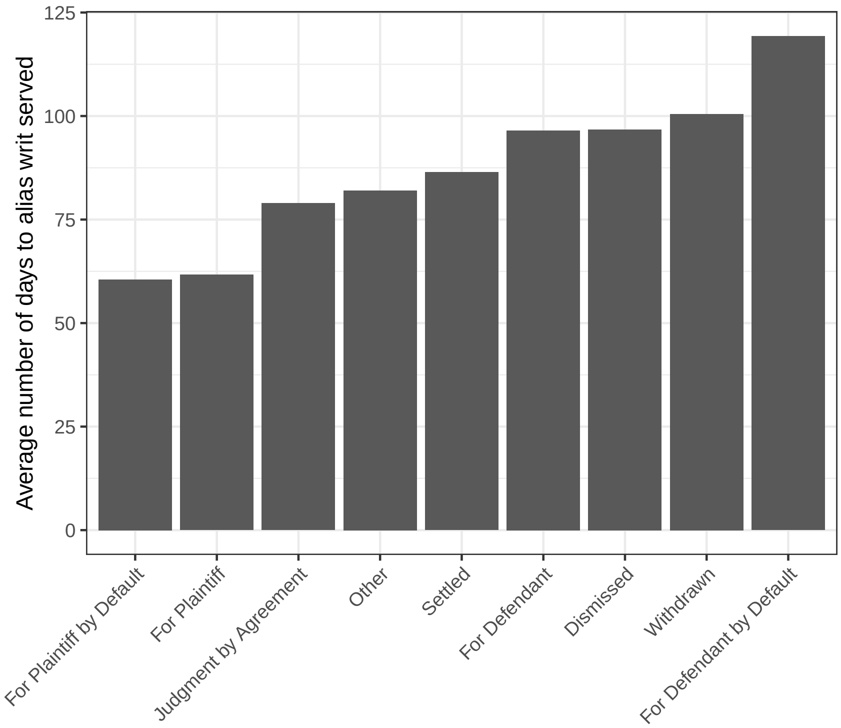

Fig. S12. Days from decision to alias writ served by judgment type: 2005-2021. 58,155 eviction proceedings with served alias writs of possession, 73,079 defendants

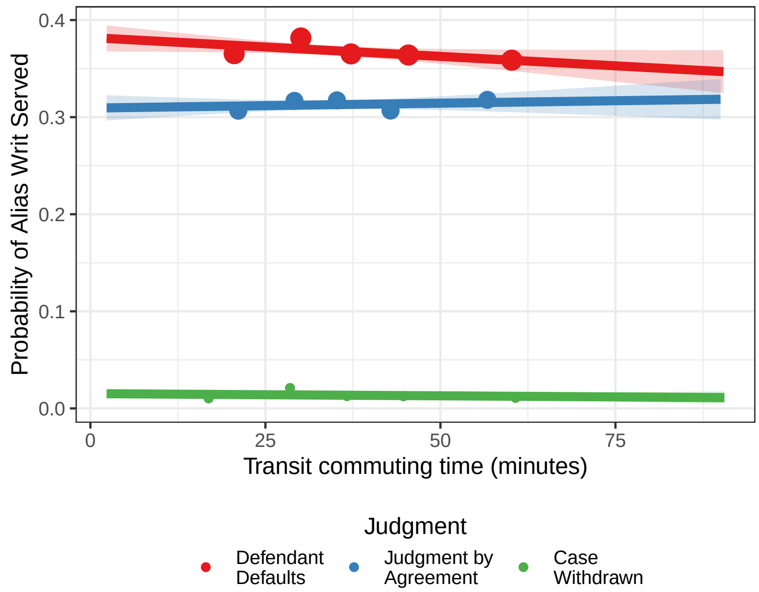

Fig. S13. Bivariate regression of alias writ served on commuting time. 53,578 unique buildings, 181,958 eviction proceedings, 232,709 defendants. Points denote binned averages of the outcome variable. Robust standard errors are clustered on building.

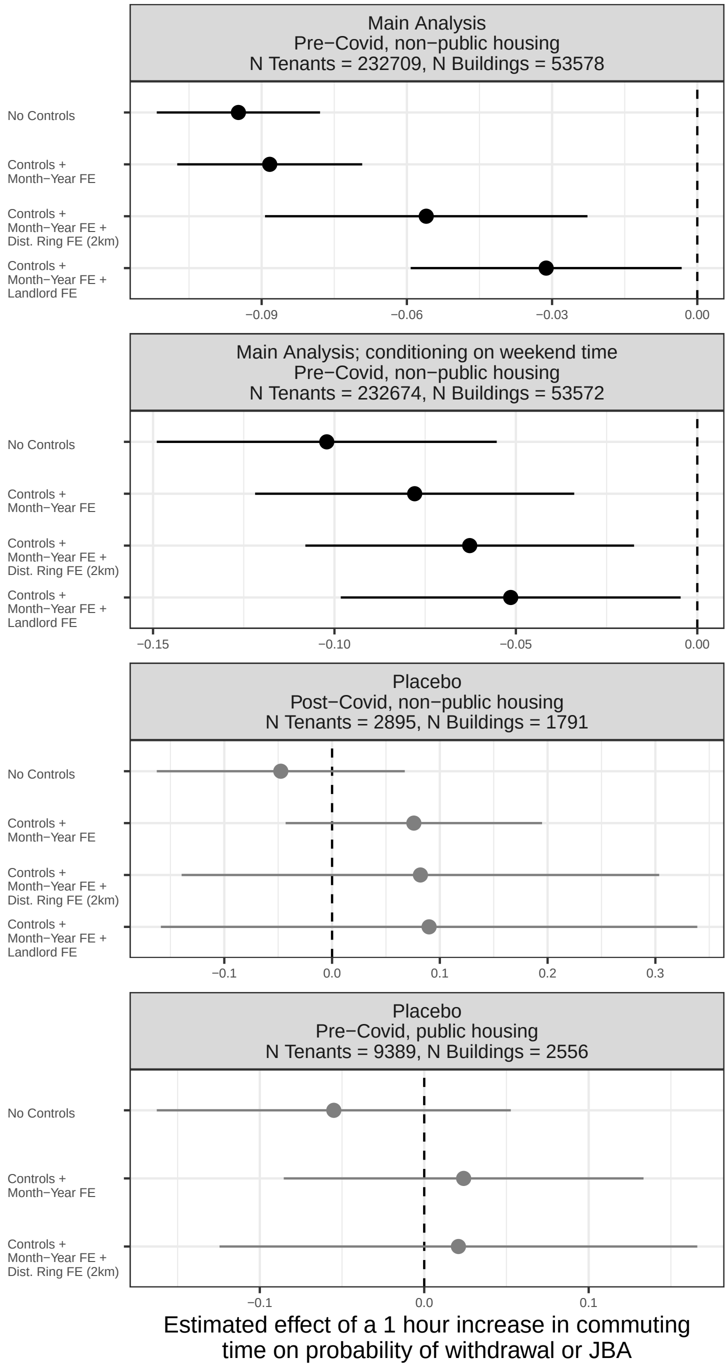

Fig. S14. Estimated average effects of a 1 hour increase in commuting time on probability of a withdrawal or a judgment by agreement. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint and whether the building is classified as an apartment. Lines denote cluster-robust 95% confidence intervals.

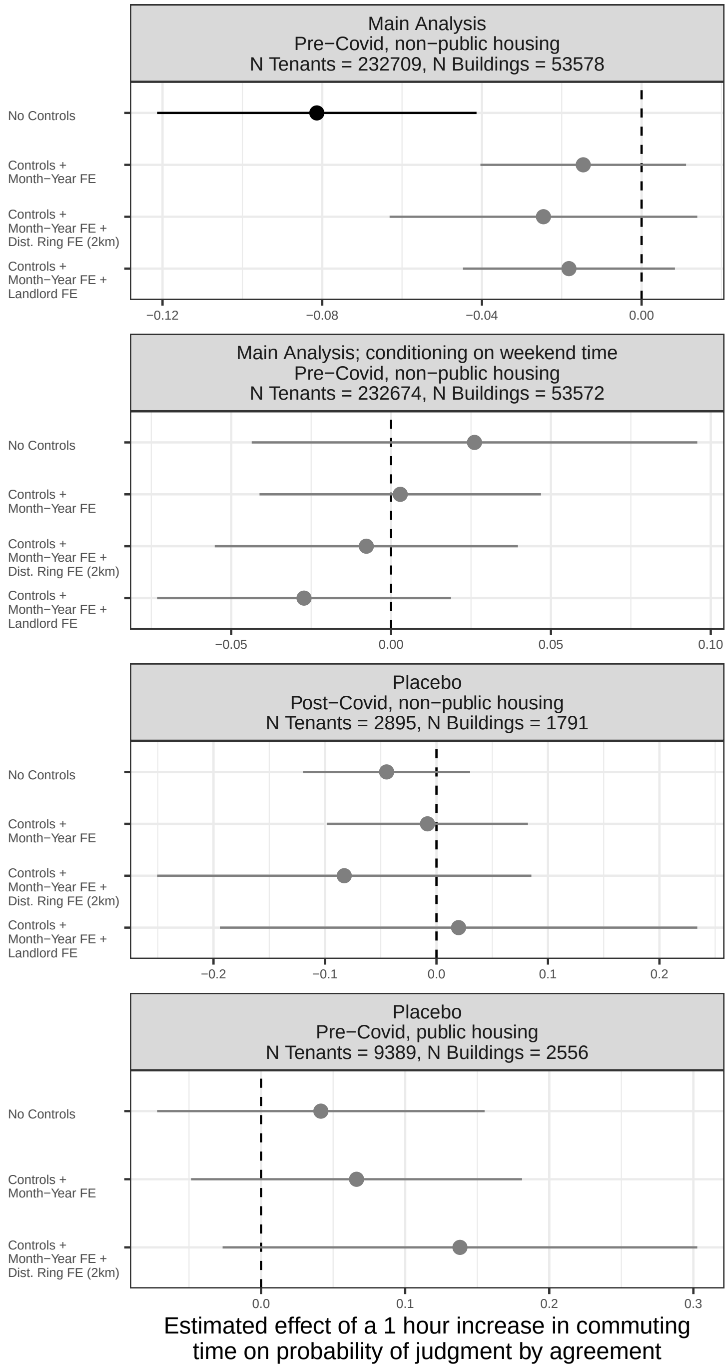

Fig. S15. Estimated average effects of a 1 hour increase in commuting time on probability of a withdrawal or a judgment by agreement. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint and whether the building is classified as an apartment. Lines denote cluster-robust 95% confidence intervals.

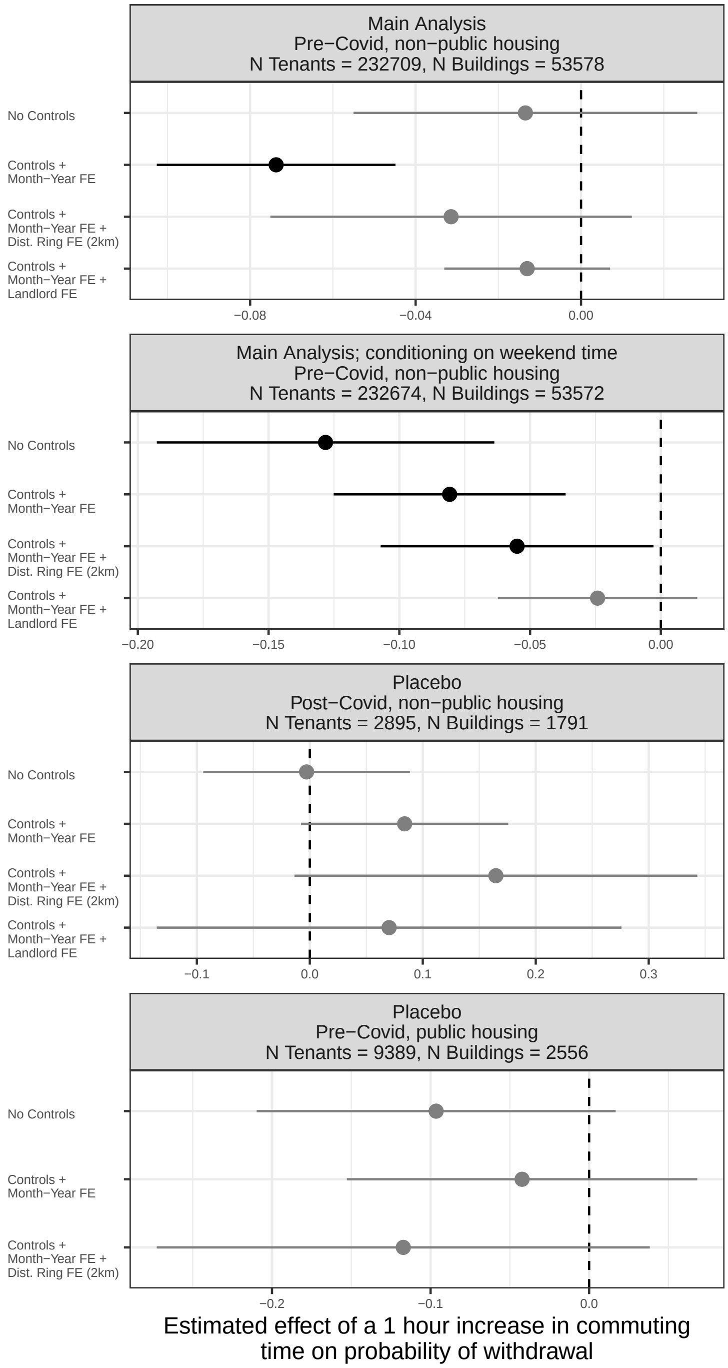

Fig. S16. Estimated average effects of a 1 hour increase in commuting time on probability of a withdrawal. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint and whether the building is classified as an apartment. Lines denote cluster-robust 95% confidence intervals.

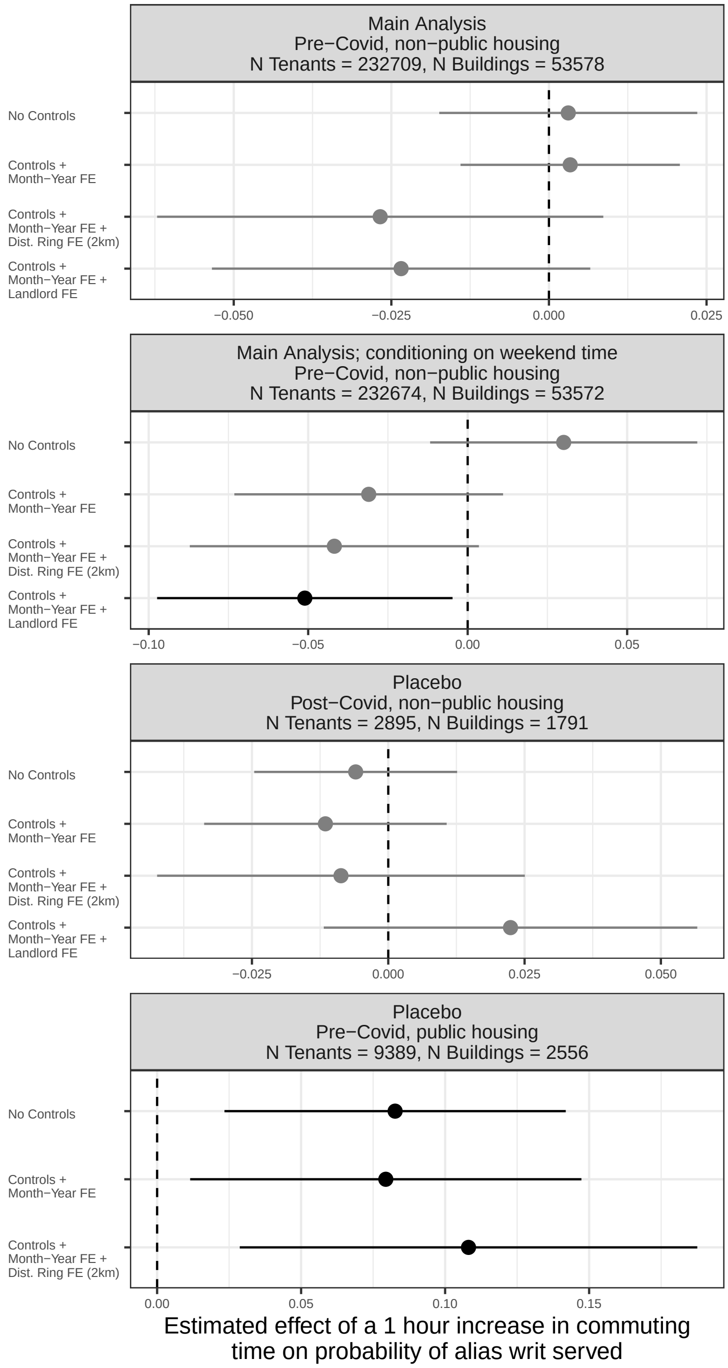

Fig. S17. Estimated average effects of a 1 hour increase in commuting time on probability of an alias writ served. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint and whether the building is classified as an apartment. Lines denote cluster-robust 95% confidence intervals.

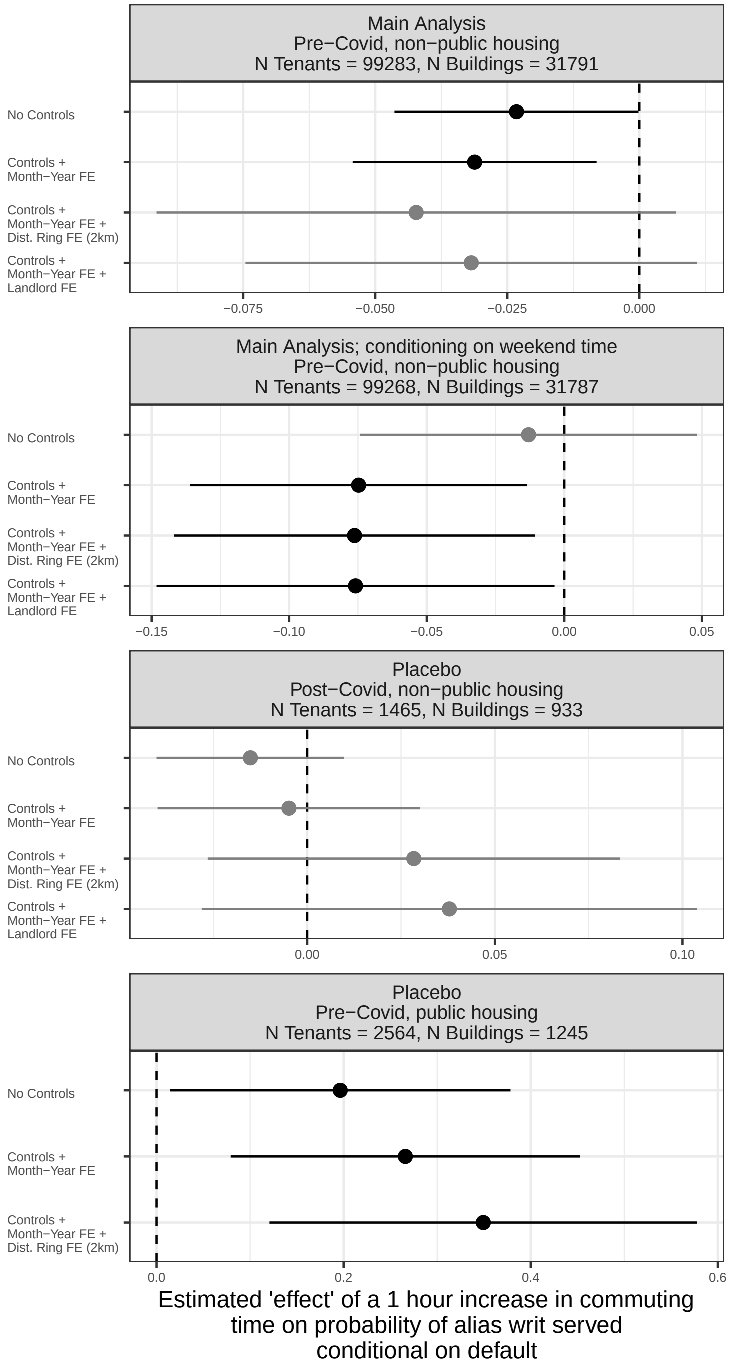

Fig. S18. Estimated average effects of a 1 hour increase in commuting time on probability of an alias writ served conditional on default judgment for plaintiff. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint and whether the building is classified as an apartment. Lines denote cluster-robust 95% confidence intervals.

7. Absence of seasonality in treatment effects

We also investigate whether the effect of commuting varies throughout the year. One potential source of heterogeneity could be the variation in commute quality by season. If busses are more likely to suffer problems in the winter than in other seasons, for example, then we might see the effect of commute time as defendants' commutes become potentially longer and more unpredictable during the colder months. We estimate our main regression, adjusting for all confounders, on subsets of our data split by three-month seasons. Figure S19 plots the point estimates and shows that there is no substantial variation across the four groups. All effect estimates remain positive, statistically significant, and roughly the same magnitude.

We also test for the interaction between treatment and indicators for each season in the pooled regression model and likewise find no statistically significant differences in the effect of commuting time on default rate between the spring, summer, fall and winter months. The results do not show any detectable seasonality in the treatment effect.

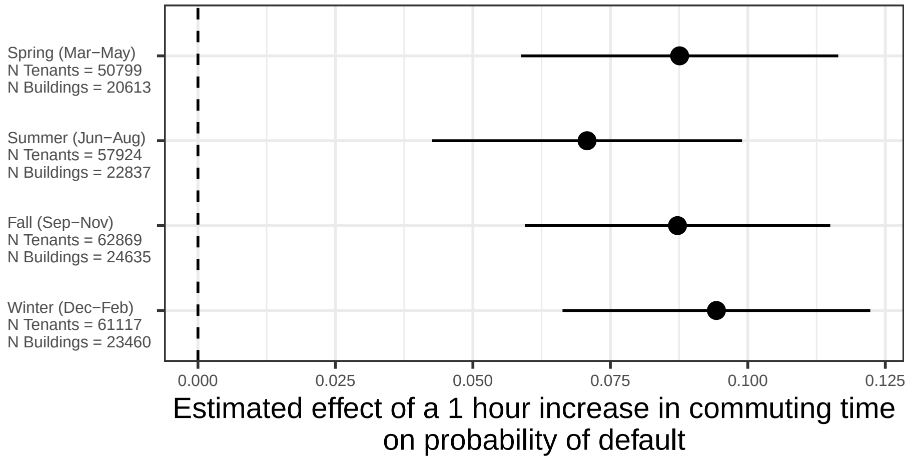

Fig. S19. Estimated average effects of a 1 hour increase in commuting time on probability of tenant default - heterogeneity by season. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, a quadratic polynomial of census block % Hispanic, estimated monthly rent from eviction complaint and whether the building is classified as an apartment. Month-year fixed effects included in all regressions. Lines denote cluster-robust 95% confidence intervals

8. Regression tables for main text results

This section provides the full regression tables for the OLS results presented in Figure 5 of the main text.

Table S1. OLS regressions of default judgment on commuting time - Main Analysis, Pre-Covid

	Model 1	Model 2	Model 3	Model 4
Transit Duration (hours)	0.1219^{**}	0.0857^{**}	0.0487^{**}	0.0382^{**}
	(0.0094)	(0.0097)	(0.0169)	(0.0145)
Median income (logged)		-0.0090	-0.0101	0.0070
		(0.0067)	(0.0068)	(0.0073)
Median contract rent (logged)		0.0191	0.0235	-0.0208
		(0.0140)	(0.0139)	(0.0137)
Block share white		0.1106^{**}	0.1008^{**}	0.0386
		(0.0333)	(0.0308)	(0.0332)
Block share white^2		-0.0925^{**}	-0.0780^{*}	-0.0114
		(0.0331)	(0.0317)	(0.0377)
Block share Hispanic		0.1050^{**}	0.0918^{**}	0.0328
		(0.0351)	(0.0334)	(0.0388)
Block share Hispanic^2		-0.1003^{*}	-0.0698	-0.0212
		(0.0425)	(0.0411)	(0.0482)
Apartment building		0.0359^{**}	0.0351^{**}	0.0128^{*}
		(0.0045)	(0.0046)	(0.0062)
Ongoing rent missing indicator		-0.0106	-0.0105	-0.0057
		(0.0075)	(0.0075)	(0.0079)
Ongoing rent $\times$ Ongoing rent non-missing		-0.0000^{*}	-0.0000^{*}	-0.0000^{**}
		(0.0000)	(0.0000)	(0.0000)
Month-Year FEs	No	Yes	Yes	Yes
Ring FEs (2km)	No	No	Yes	No
Landlord FEs	No	No	No	Yes
R^2	0.0033	0.0106	0.0112	0.1690
Num. obs.	232709	232709	232709	232709
N Clusters (building)	53578	53578	53578	53578

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

Notes: Estimates from a linear probability model estimated by ordinary least squares. Cluster-robust standard errors in parentheses. Intercept and fixed effects parameters omitted to conserve space.

Table S2. OLS regressions of default judgment on commuting time - Main Analysis, Pre-Covid, Adjusting for weekend commute

	Model 1	Model 2	Model 3	Model 4
Transit Duration (hours)	0.0892^{**}	0.0658^{**}	0.0487^{*}	0.0589^{*}
	(0.0269)	(0.0230)	(0.0233)	(0.0238)
Weekend transit duration (hours)	0.0348	0.0214	-0.0001	-0.0247
	(0.0285)	(0.0244)	(0.0272)	(0.0248)
Median income (logged)		-0.0095	-0.0101	0.0081
		(0.0067)	(0.0069)	(0.0074)
Median contract rent (logged)		0.0188	0.0235	-0.0214
		(0.0138)	(0.0137)	(0.0137)
Block share white		0.1103^{**}	0.1008^{**}	0.0380
		(0.0332)	(0.0308)	(0.0333)
Block share white^2		-0.0911^{**}	-0.0780^{*}	-0.0120
		(0.0332)	(0.0317)	(0.0378)
Block share Hispanic		0.1067^{**}	0.0917^{**}	0.0303
		(0.0352)	(0.0337)	(0.0390)
Block share Hispanic^2		-0.1027^{*}	-0.0695	-0.0174
		(0.0427)	(0.0417)	(0.0485)
Apartment building		0.0356^{**}	0.0351^{**}	0.0131^{*}
		(0.0045)	(0.0045)	(0.0062)
Ongoing rent missing indicator		-0.0107	-0.0105	-0.0057
		(0.0075)	(0.0075)	(0.0079)
Ongoing rent $\times$ Ongoing rent non-missing		-0.0000^{*}	-0.0000^{*}	-0.0000^{**}
		(0.0000)	(0.0000)	(0.0000)
Month-Year FEs	No	Yes	Yes	Yes
Ring FEs (2km)	No	No	Yes	No
Landlord FEs	No	No	No	Yes
R^2	0.0033	0.0106	0.0112	0.1691
Num. obs.	232674	232674	232674	232674
N Clusters (building)	53572	53572	53572	53572

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

Notes: Estimates from a linear probability model estimated by ordinary least squares. Cluster-robust standard errors in parentheses. Intercept and fixed effects parameters omitted to conserve space.

Table S3. OLS regressions of default judgment on commuting time - Placebo analysis, Post-Covid

	Model 1	Model 2	Model 3	Model 4
Transit Duration (hours)	0.1152	−0.0269	−0.0518	−0.0970
	(0.0605)	(0.0606)	(0.1118)	(0.1258)
Median income (logged)		−0.0555	−0.0410	−0.1402*
		(0.0411)	(0.0430)	(0.0667)
Median contract rent (logged)		0.1133	0.1087	0.1770
		(0.0748)	(0.0791)	(0.1284)
Block share white		0.1453	0.0887	0.3653
		(0.1725)	(0.1729)	(0.2665)
Block share white^2		−0.0721	−0.0320	−0.0864
		(0.1901)	(0.1864)	(0.3218)
Block share Hispanic		0.1150	0.1815	0.1683
		(0.1981)	(0.2036)	(0.3612)
Block share Hispanic^2		−0.2128	−0.2752	−0.2583
		(0.2603)	(0.2716)	(0.4480)
Apartment building		0.1018**	0.1034**	0.0617
		(0.0261)	(0.0275)	(0.0603)
Ongoing rent missing indicator		0.3399**	0.3418**	0.0700
		(0.1060)	(0.1080)	(0.2306)
Ongoing rent × Ongoing rent non-missing		0.0000	0.0000	−0.0000
		(0.0000)	(0.0000)	(0.0001)
Month-Year FEs	No	Yes	Yes	Yes
Ring FEs (2km)	No	No	Yes	No
Landlord FEs	No	No	No	Yes
R^2	0.0028	0.1046	0.1092	0.6521
Num. obs.	2895	2895	2895	2895
N Clusters (building)	1791	1791	1791	1791

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

Notes: Estimates from a linear probability model estimated by ordinary least squares. Cluster-robust standard errors in parentheses. Intercept and fixed effects parameters omitted to conserve space.

Table S4. OLS regressions of default judgment on commuting time - Placebo analysis, Public housing

	Model 1	Model 2	Model 3
Transit Duration (hours)	0.0412	-0.0011	-0.0291
	(0.0548)	(0.0561)	(0.0738)
Median income (logged)		-0.0148	-0.0041
		(0.0189)	(0.0201)
Median contract rent (logged)		0.0046	0.0122
		(0.0323)	(0.0326)
Block share white		-0.0647	-0.0452
		(0.1311)	(0.1316)
Block share white^2		0.0711	0.0238
		(0.2209)	(0.2307)
Block share Hispanic		0.0038	0.0186
		(0.0998)	(0.1028)
Block share Hispanic^2		0.0363	0.0189
		(0.1250)	(0.1297)
Apartment building		0.0175	0.0235
		(0.0153)	(0.0170)
Ongoing rent missing indicator		-0.0583^{**}	-0.0565^{**}
		(0.0170)	(0.0170)
Ongoing rent $\times$ Ongoing rent non-missing		-0.0001^{**}	-0.0001^{**}
		(0.0000)	(0.0000)
Month-Year FEs	No	Yes	Yes
Ring FEs (2km)	No	No	Yes
R^2	0.0001	0.0373	0.0381
Num. obs.	9389	9389	9389
N Clusters (building)	2556	2556	2556

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

Notes: Estimates from a linear probability model estimated by ordinary least squares. Cluster-robust standard errors in parentheses. Intercept and fixed effects parameters omitted to conserve space.

9. Analysis of evictions in Harris County, TX

To evaluate whether our core finding generalizes beyond the specific context of Philadelphia, we examine another dataset of evictions from Harris County, Texas which contains the city of Houston. There are notable differences between eviction procedures in Harris County and Philadelphia which we describe below.

We obtained from EvictionLab a dataset of 971,633 eviction proceedings filed in Harris County between 2000 and 2018. We removed all entries that had a missing filing date, that did not contain a valid Harris County address for the defendant, or that did not contain data on the courthouse where the proceeding was filed. We further removed any non-residential evictions as well as those where the nature of the claim was clearly marked as not an eviction case. We geocoded each valid address to a latitude and longitude using Google's Geocoding API and dropped all proceedings that did not fall within a census tract in Harris County. This process left a total of 933,158 proceedings.

Harris County civil cases are assigned to one of sixteen Justice of the Peace courts organized into eight precincts with two courts per precinct. These precincts are geographically bounded and the boundaries have remained relatively constant over the 2000-2018 period. Likewise, the addresses of the courts have only changed during this period for two of the courthouses and these location changes have been negligible (moving at most two or three blocks) resulting in 18 unique courthouses buildings in our dataset. The vast majority of cases in our data are correctly assigned to one of the two courthouses in their geocoded precinct. However, we do observe a small proportion that are filed in courthouses that are not in the correct precinct. We find that many of these appear to be dismissed and believe that they are either errors in the data or misfiled cases. We drop these cases as well, leaving a total of 910,963 evictions. 2,356 of these are public housing evictions involving a local Housing Authority as a plaintiff.

For each of these evictions, we obtained the commuting distance and time from the property to its assigned courthouse using Google's Distance Matrix API. Because public transit is much less heavily used in Houston and the surrounding area and many of the courthouses are entirely inaccessible by transit, we use the driving time estimates instead. While the vast majority of addresses were correctly processed by the Distance Matrix API, a handful returned errors or zero results. We drop these from our data leaving 910,225 evictions of which 2,354 involve public housing.

We obtain the outcome of each eviction proceeding using the classification assigned by the Office of Court Administration. Only 1,353 cases are missing a classification, either because of data entry errors or because the eviction proceeding has not been completed. We also remove these cases from our data. In total, our primary dataset contains a total of 908,871 eviction proceedings of which 906,520 involve non-public housing cases.

As in Philadelphia, the modal outcome of an eviction case is a default judgment against the tenant. We understand that there is effectively no reopening procedure at the justice of the peace level in Harris County: defaulted tenants may file a new civil case in state court, depositing a month's rent in advance and an additional sum every month the case is pending. That is, defaults are terminal events in Harris County even more than they are in Philadelphia. Figure S20 plots the share of each outcome among the cases in our dataset. About 37 percent of cases are defaults in favor of the plaintiff. The next most common outcome (31 percent) is a dismissal, either by the plaintiff withdrawing the case or by the judge. In contrast to Philadelphia, settlements or judgments by agreement are much rarer and many more cases are recorded as having been heard by a judge. 17% of all cases result in the plaintiff winning at trial. It is possible–and we in fact suspect—that some of these "victories" are actually negotiated settlements classified as judgments.

We replicate our Philadelphia design in Harris County and estimate the average effect of commuting time via driving on the probability of default. First, we subset our data down to only non-public housing evictions. Next, merge our dataset to the same set of ACS and Census tract/block-level covariates that we use in our Philadelphia analysis. While nearly all observations are geolocated to census tracts with non-missing data on median income and median contract rent, a larger share — about 9.8 percent — are geolocated to census blocks that are recorded as having zero residents in 2010. Since Harris County encompasses areas outside of Houston that can be much less dense than the city core, this is expected, but unfortunately means that we cannot obtain data on racial demographics in these blocks. After dropping all observations with missing tract or block covariate data, our dataset consists of 815,365 evictions across 125,335 unique properties.

Figure S21 plots the simple bivariate regression of default on driving commutes and finds a strong positive relationship between the two in our dataset. Because the sizes of the Justice of the Peace precincts differ substantially we might be concerned that there are some courthouses that are generally more difficult to reach and that this might be correlated with default rates. It is the case that some courthouses have default rates that significantly exceed the Harris County average and others that are consistently below average. This may also result from very differing procedures between courthouses about how and whether to contact missing tenants. However, we find that even after adjusting for courthouse fixed effects — comparing commuting times for cases assigned to the same courthouse — we find the same positive relationship persists on average. Going from a commute that is 10 minutes below the courthouse mean to one that is 10 minutes above corresponds to a roughly 6-7 percentage point increase in the probability of default.

We attempt to further address possible unobserved confounding by adjusting for observed covariates at the census tract and block level. As in the main results, we control for the log of median income, logged median block contract rent, a quadratic polynomial of the share of white residents in the census block and a quadratic polynomial of the share of hispanic residents in the census block. We also include courthouse by filing month-year fixed effects (18 courthouses; 228 unique month-year combinations) to address confounding due to differences across the underlying courthouses and due to potential time trends in defaults and filing patterns. Figure S22 plots the estimated average effect of a 10 minute increase in driving time on the probability of default. [‡] We find that even after adjusting for covariates, we obtain a strong positive treatment effect estimate. Increasing a tenant's commuting time by 10 minutes raises their probability of default, on average, by about 3.3 percentage points. Notably, the covariates do very little to change the point estimate — the adjusted and unadjusted results are remarkably similar to one another.

Due to limitations with this dataset, we lack information on rental prices: we can only measure rent at the census tract level. We also lack information about building type and do not have data that permits us to link landlord LLCs. Also, because driving time is much more highly collinear with direct distance, there is very little residual variation in commutes if we condition on the direct measured distance to the courthouse. Because of this, we are unable to apply either the landlord fixed effects, weekend effects or the ring fixed effects strategies from the main analysis. Due to the dearth of public housing cases in these data and the absence of any information about a different treatment of those evictions, can not public housing as a control. And since our data pre-date 2020, we can not use remote access as a placebo test (and it at best unclear if the justices of the peace all permitted remote access even when mandated to do so).

However, the presence of multiple courthouses in a precinct permits another novel identification strategy based on the fact that assignment to a courthouse is non-deterministic. Two evictions occurring in the same month on the same block or even in the same property could be assigned to entirely different courthouses. We therefore implement another analysis that compares default rates in a given month *across* courthouses for defendants that live on the same census block (a total of 395,964 unique fixed effect parameters). We do an even more extreme version of this analysis that looks at defendants living in the same *property* (455,045 unique fixed effects). These designs rule out any confounding due to unobserved geographic factors and the latter analysis even addresses concerns about unobserved landlord selection effects by narrowly comparing tenants being evicted from the same exact building. In all of our analyses (Figure S22, Table S5), we obtain very similar point estimates: a 10 minute increase in commuting time corresponds to about a 3-4 percentage point increase in the probability of a default judgment.

Overall our results from Harris County strengthen the generalizability of our primary findings from Philadelphia. Even in a jurisdiction with radically different operating procedures and transportation options, we find a similar pattern linking commute time and distance to tenant defaults in eviction proceedings. While we are unable to implement the exact same identification strategies as in Philadelphia due to data availability, we were able to implement one novel strategy that was impossible for Philadelphia — leveraging the presence of multiple courthouses in Harris County. When we compare neighboring evictions assigned to different courthouses, we still find a consistent and strong relationship between the length of the tenant's commute and the probability that they default. Across all analyses we find consistent evidence supporting this core hypothesis.

[‡] Because the typical range of driving times is much more compressed than transit times, we provide estimates for a 10 minute increase rather than a one hour increase as they are more representative of the kind of variation we see within a Justice of the Peace precinct.

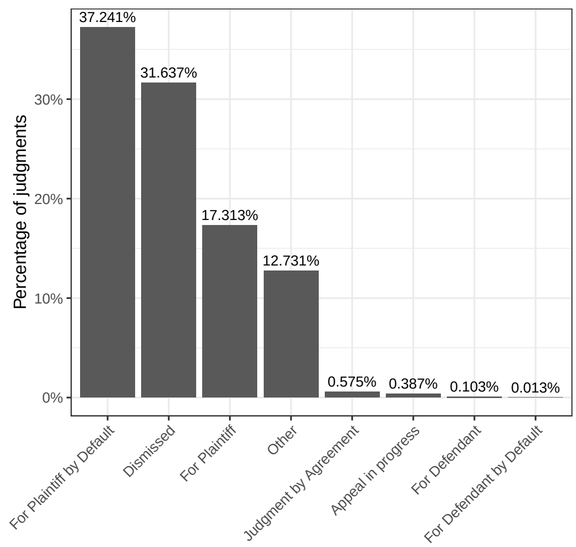

Fig. S20. Distribution of outcomes in Harris County, TX Eviction proceedings: 2000-2018. 908,871 eviction proceedings

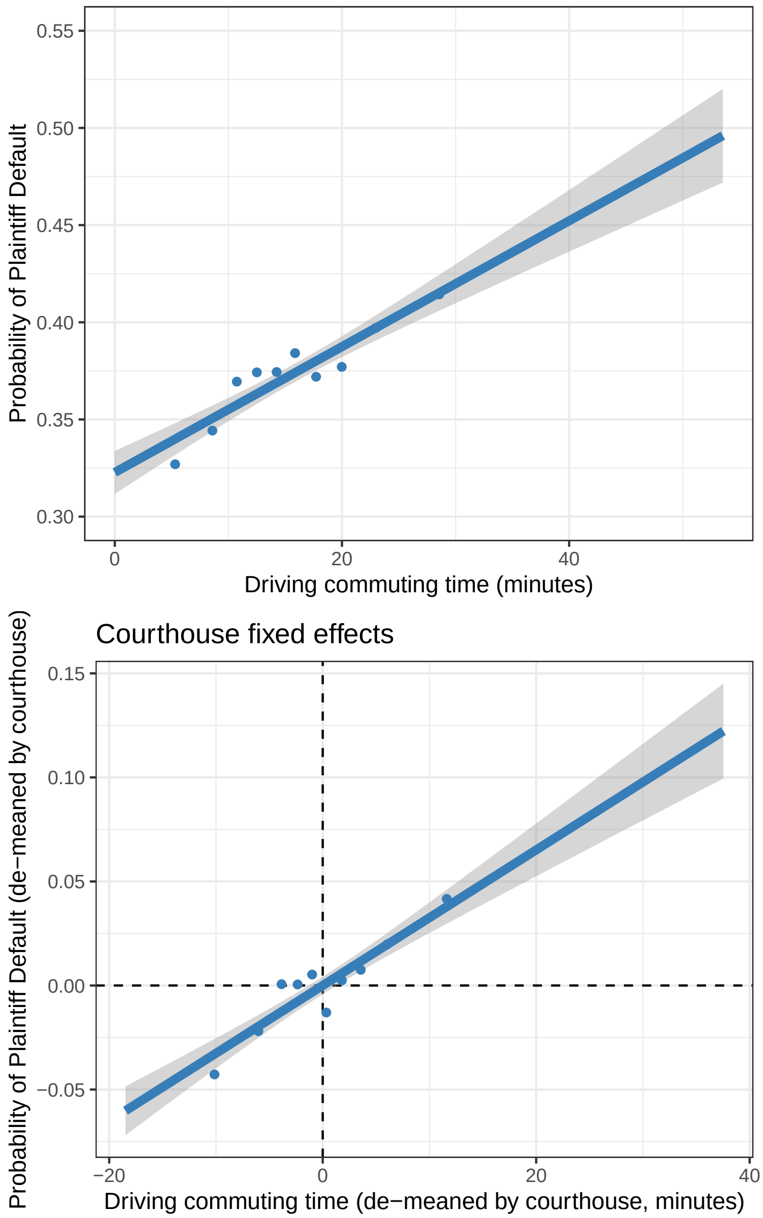

Fig. S21. Bivariate regression of default probability on driving time - Harris County. 125,335 unique buildings, 815,365 eviction proceedings. Points denote binned averages of the outcome variable. Robust standard errors are clustered on building. Second plot de-means commuting times and default rates by courthouse (18 unique courthouses)

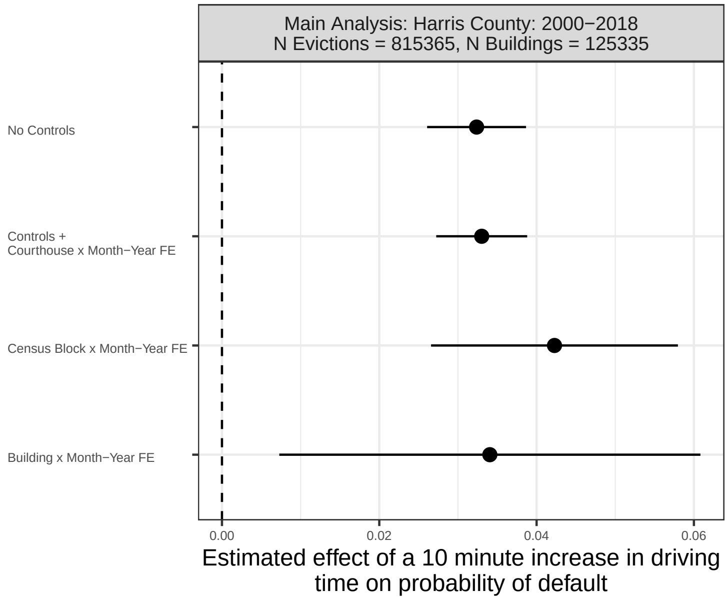

Fig. S22. Estimated average effects of a 10 minute increase in driving time on probability of tenant default - Harris County. Estimates from a linear probability model estimated by ordinary least squares. Covariates include census tract median income (logged), census tract median contract rent (logged), a quadratic polynomial of census block % white, and a quadratic polynomial of census block % Hispanic. Lines denote cluster-robust 95% confidence intervals clustered at the building level.

Table S5. OLS regressions of default judgment on commuting time - Harris County (2000-2018)

	Model 1	Model 2	Model 3	Model 4
Driving duration (tens of minutes)	0.0324^{**}	0.0330^{**}	0.0423^{**}	0.0341^{*}
	(0.0032)	(0.0030)	(0.0080)	(0.0137)
Median income (logged)		0.0080		
		(0.0085)		
Median contract rent (logged)		-0.1041^{**}		
		(0.0120)		
Block share white		-0.2354^{**}		
		(0.0405)		
Block share white^2		0.2351^{**}		
		(0.0390)		
Block share Hispanic		0.1858^{**}		
		(0.0314)		
Block share Hispanic^2		-0.1062^{**}		
		(0.0294)		
Courthouse × Month-Year FE	No	Yes	Yes	Yes
Census Block × Month-Year FE	No	No	Yes	No
Building × Month-Year FE	No	No	No	Yes
R^2	0.0020	0.0381	0.5761	0.6436
Num. obs.	815365	815365	815365	815365
N Clusters (building)	125335	125335	125335	125335

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

Notes: Estimates from a linear probability model estimated by ordinary least squares. Cluster-robust standard errors in parentheses. Intercept and fixed effects parameters omitted to conserve space.

References

1. Cinelli C, Hazlett C (2020) Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82(1):39–67.
2. Eisenhard S, et al. (2020) Reducing default judgments in philadelphia's landlord-tenant court, (Sheller Center for Social Justice), Technical report.
3. Hoffman DA, Strezhnev A (2022) Leases as forms. *Journal of Empirical Legal Studies* 19(1).
4. Summers N (2021) Civil probation. *Stanford Law Review* 75.
5. Reinvestment Fund (2020) Resolving landlord-tenant disputes: An analysis of judgments by agreement in philadelphia's eviction process. *Reinvestment Fund Policy Brief.*
6. Seron C, Frankel M, Ryzin GV, Kovath J (2001) The impact of legal counsel on outcomes for poor tenants in new york city's housing court: Results of a randomized experiment. *Law & Society Review* 35(2):419–434.
7. Elwert F, Winship C (2014) Endogenous selection bias: The problem of conditioning on a collider variable. *Annual review of sociology* 40:31.