MetroSight
Economics. Research. Clarity. Insight.

Rethinking rent

Flexible rent payments can improve outcomes for property owners while helping residents manage cash flow.

Daniel Shoag, Ph.D.
Issi Romem, Ph.D.

June 2026

Executive Summary

Policy and industry discussions about housing stability tend to focus on affordability. Yet one important and often overlooked driver of late rent and renter instability is the mismatch between household cash flow and fixed rent payment deadlines. When household financial shocks hit as rent comes due, this timing misalignment can trigger a destabilizing spiral that affects both renters and property owners.

A substantial share of U.S. households lack the financial buffers to absorb even modest disruption. An illness or medical bill, car repairs, or a late paycheck can quickly lead to late or missed rent payments. These disruptions can compound over time, increasing the risk of housing instability and imposing significant financial and psychological costs on renters.

Tenant turnover also imposes meaningful costs on multifamily operators. It's not just the loss of rental income during unit vacancies. Replacing a resident can require thousands of dollars in repairs, marketing, leasing incentives, and administrative time, costs that can materially affect operating performance.

Taken together, data on household financial stability and multifamily operating costs point to the need for solutions that stabilize rent payments for property operators while alleviating financial stress for residents.

Several services now offer residents a way to split rent payments, pay in installments, or defer payments to better match income schedules, effectively resolving a structural timing mismatch between when income arrives and when rent is due and heading off the cascade of negative outcomes that can follow a missed rent payment.

This report examines the effects of one such service, Flex, which facilitates on-time rent payments to property owners while allowing residents to split their rent into multiple payments.[1]

Using data from financial provider Flexible Finance Inc., we find that the availability of Flex improved rent payment reliability, corresponded to longer resident tenure and lower vacancy, and coincided with stronger property operating performance.

Key findings

- **Flex improved rent-payment timing.** Properties offering Flex had higher current-payment rates and lower short-term delinquency. Across the main estimates, on-time rents were about 3 to 5 percentage points higher. We do not find a statistically significant effect on rents more than 30 days late, consistent with Flex addressing timing rather than inability to pay.
- **Flex was associated with greater resident stability.** Properties offering Flex showed longer resident tenure and lower vacancy in adjusted comparisons. These outcomes are farther

[1] Flex's product is structured around a bank-issued line of credit regulated under the Truth in Lending Act, Regulation Z (https://getflex.com/policies). The product charges no late fees and accrues no compounding interest. The line of credit must be repaid in full before it can be used again, meaning balances cannot accumulate across billing cycles and renters cannot layer new credit on top of outstanding debt. Because the findings in this report are derived exclusively from this specific product structure, their external validity might not extend to similar products with different credit terms or fee designs, which might produce materially different financial outcomes for renters.

downstream than rent-payment timing, but they are consistent with the same mechanism: fewer payment disruptions reducing resident instability and vacancy pressure.

- **Flex coincided with stronger operating outcomes.** Flex properties showed higher NOI and lower turnover-related, collection-related, and concession costs in the reported comparisons. The NOI evidence is positive but imprecise, so it should be read as directional evidence of stronger operating performance rather than as a precise estimate of the effect's magnitude.

The main patterns are broadly consistent across several observational approaches, with the strongest support for rent-payment timing. The results underscore why rent timing matters: when a temporary cash-flow mismatch becomes a late rent payment, the consequences can extend beyond the current month's balance to resident stability, vacancy, collections, and property operations.

Introduction

Most residential leases require residents to pay their rent once each month, typically on the first day or soon thereafter. While specific terms can vary by jurisdiction and property, this fundamental practice in the United States is nearly universal.

This practice creates a structural mismatch between when rent is due and when income is received. Approximately 56 percent of U.S. private-sector workers are paid bi-weekly. Another 19 percent are paid once or twice a month.[2]

For most renters, that means their obligations don't align neatly with their income, a mismatch that can have financial implications for residents and property managers alike.

For residents living paycheck to paycheck, the gap between when money arrives and when rent is due can create significant financial strain. A delayed paycheck can result in late fees, which averaged approximately $85 as of late 2024, or foregone spending on essentials.[3]

This misalignment becomes especially consequential in the face of adverse financial shocks. Events such as job loss, reduced hours, medical emergencies, or unexpected car repairs can abruptly disrupt cash flow. Because rent is both large and inflexible—due in full on a fixed date—these shocks often leave residents with few options.

Even short-term income disruptions can cascade into late fees, mounting debt, or difficult trade-offs. For households with little or no savings, a rigid monthly rent schedule amplifies the financial impact of otherwise temporary and often manageable setbacks.

---

[2] Hayes M, Northup J. *More paydays, more fairness? How pay frequency shapes worker perceptions of compensation.* ADP Research Institute. Published June 25, 2025. Accessed March 7, 2026. https://www.adpresearch.com/more-paydays-more-fairness-how-pay-frequency-shapes-worker-perceptions-of-compensation/

[3] Consumer Financial Protection Bureau. *Behind on Rent? Examining Rental Housing Delinquencies in New Payment Data.* Published January 24, 2025. Accessed April 7, 2026. https://www.consumerfinance.gov/data-research/research-reports/behind-on-rent-examining-rental-housing-delinquencies-in-new-payment-data/

The risk is especially pronounced for cost-burdened renters. More than 18.1 million U.S. households spend at least 35 percent of their gross income on rent, and 37 percent of adults say they would be unable to cover a $400 emergency expense with cash.[4,5]

When cash flow disruptions occur, they compound financial stress, increasing the likelihood of delinquency and even eviction.

For property owners and operators, late or missed rents don't just disrupt cash flow, they compound risk. Missed payments can cascade into vacancies, with the cost of turnover—including lost rent, marketing, resident screening, cleaning, and repair—ranging from $1,500 to $3,500 per move-out.[6]

The timing and frequency of payment disruptions can be difficult to anticipate, undermining financial forecasting and leading to distorted property-level cash flow projections. Collectively, these effects can erode net operating income and weaken property performance.[7]

The scale of these issues is substantial. More than one-third of U.S. households rent their homes.[8] Of these, between 14 percent and 23 percent reported being behind on their rent payments between 2021 and 2024.[9] These figures speak to the importance of identifying ways to improve rent payment reliability for property operators while reducing financial stress for residents.

In response to this problem, financial companies are offering services that allow renters to split their monthly rent payments into more frequent, smaller remittances. One of these, Flexible Finance Inc., or Flex, enables residents to split their monthly rent into two scheduled payments to improve cash flow and avoid late charges. For a monthly fee, Flex users can pay a portion of their rent–half of it or less–at the beginning of the month. Flex advances the remaining amount to the property owner, and residents pay Flex the remaining balance by the end of the month.[10]

By converting a single large financial obligation into smaller payments that better align with income flows, such services can help renters manage cash flow and reduce the risk of late fees, overdrafts, or eviction.

This study evaluates the relationship between the availability of the Flex service and property- and unit-level outcomes.

---

[4] U.S. Census Bureau. *American Community Survey 1-Year Estimates: Selected Housing Characteristics (Table DP04).* Accessed March 8, 2026. https://data.census.gov. More than 18.1 million U.S. households spend 35 percent or more of their gross income on rent.

[5] Federal Reserve Board. *Economic Well-Being of U.S. Households in 2024.* Washington, DC: Federal Reserve Board; 2025. Accessed March 8, 2026. https://www.federalreserve.gov/consumerscommunities/shed.htm

[6] National Apartment Association. *The resident search journey.* Published June 2024. Accessed March 7, 2026. https://www.naahq.org/sites/default/files/2024-06/NAA_Resident%20Journey_Final.pdf

[7] Romem I, Shoag D. *Behind the high cost of rent.* MetroSight. Published February 25, 2025. Accessed March 7, 2026. https://www.metrosight.com/articles/behind-the-high-cost-of-rent

[8] U.S. Census Bureau. *More renter households are cost burdened, especially among Black and Hispanic renters.* Published September 2024. Accessed March 7, 2026. https://www.census.gov/newsroom/press-releases/2024/renter-households-cost-burdened-race.html

[9] Consumer Financial Protection Bureau. *Behind on rent? Examining rental housing delinquencies in new payment data.* Published January 24, 2025. Accessed March 7, 2026. https://www.consumerfinance.gov/data-research/research-reports/behind-on-rent-examining-rental-housing-delinquencies-in-new-payment-data/

[10] As of March 2026, the monthly fee charged by Flex was $14.99, plus additional property and transaction fees, where applicable. For more information, see the company's website at https://help.getflex.com/hc/en-us/articles/360034246933-How-much-does-Flex-Rent-cost

We combine proprietary property, unit, and resident data from Flex with operational data. These data sets include information on resident aged receivables, property income statements, occupancy, unit mix, rent rolls, concessions, property metadata, and ZIP-code-level socioeconomic controls derived from the Census Bureau's American Community Survey.

We observe whether Flex is being offered at a property. Our cross-sectional dataset includes 488 properties comprising about 75,000 housing units, with underlying operational data running from July 2023 through late 2025, depending on the property and the specific data element. For most properties, the length of fully-observed history is less than a year, which all but precludes approaches that compare outcomes for the same properties over time.

We focus on the effects of offering Flex at the property level, rather than on individual subscribers. Residents who enroll are more likely to be financially strained, therefore comparing users to non-users would confound the effect of Flex with underlying differences in resident financial circumstances. Properties don't offer Flex at random either, but differences across properties are easier to account for than those between individual residents.

We estimate the effects of offering Flex on key outcomes using multiple complementary econometric strategies to compare outcomes at properties that offered the service and those that did not.

The control analysis: To ensure like-to-like comparisons, we first controlled for a rich set of characteristics, including unit count, leased rent, ZIP-level median household income, home values, vacancy rates, unemployment rates, average unit size, and bedroom count. Additionally, we used statistical methods to control for state-level characteristics that can't be readily observed, such as prevailing economic conditions.

The matching analysis: To further assess the robustness of these results, we applied a like-for-like comparison by matching each property offering Flex with the most similar property that did not. We measured similarity using a combination of the control factors applied in the first approach. These results were directionally consistent with those of the control approach, but tended to be smaller in magnitude.

Both approaches compared outcomes at properties that offered Flex with ones that did not. The control approach was informed by a broader data set that included non-Flex properties that were less comparable to properties offering Flex. This approach accounted for that by controlling for property characteristics.

In contrast, the matching approach used a more restricted sample that limited the comparison to non-Flex properties that were most similar to those offering Flex. Comparability was improved, but at the cost of leaving out potentially informative observations.

We took similar approaches in our comparison of outcomes at the unit level, again controlling for several characteristics, and also matching units at properties offering Flex with the most similar units at properties that did not. Unit-level characteristics used for both the control and matching approaches included unit rent, square footage, bedroom and bathroom counts, month-to-month lease status, property size, and the same ZIP-level demographic and economic characteristics we applied at the property level. As with the property-level comparison, we also controlled for unobservable state-level characteristics. Our unit-level results aligned with property-level results.

Throughout the report, we present the controlled and matched estimates separately. The controlled estimate uses the broader sample and adjusts for observed property, market, and state-level differences. The matched estimate uses a narrower comparison set of more similar non-Flex properties. Where estimates are shown in tables, confidence intervals are reported for both approaches so readers can see both statistical uncertainty and differences across empirical approaches.

A key limitation of this analysis is that it relies on a cross-sectional snapshot in time rather than taking a longer view that includes long observations before and after Flex adoption. A cross-sectional analysis over time, known as a panel study, has the potential to provide stronger causal evidence, but the currently available data is too limited in duration and unbalanced across properties to support this approach. As more data becomes available, it will provide an important avenue for future research.

Another key limitation is that properties that adopted Flex may have differed from those that did not in ways that are difficult to fully account for. Among the subset of Flex properties that could be observed near or prior to rollout, vacancy rates and late-rent shares tended to be somewhat lower than among the control properties at that time, though some measures of financial performance appeared weaker. Although most of these pre-treatment differences were not statistically significant, they nonetheless raise the possibility that underlying differences between properties contributed to the estimated effects along with the effect of Flex itself.

These concerns motivate the study's reliance on observable controls, matching methods, unit-level analysis and multiple sensitivity and robustness checks. Both the control and matching approaches substantially improve the comparability of Flex and non-Flex properties, but they cannot fully eliminate the possibility that unobserved differences between properties contributed to the estimated effects. It remains possible that properties adopting Flex were already on stronger operational trajectories than the control group prior to rollout. It is also possible that differences in management quality or operator sophistication contributed both to the decision to adopt Flex and to stronger operating outcomes independent of the service itself. Our sensitivity analysis suggests that unobserved factors would need to be substantially more important than the observed differences already accounted for in the study—including differences in property size, rent levels, local market conditions, vacancy rates, and neighborhood income—in order to fully explain away the estimated effects. Still, some uncertainty around causal interpretation necessarily remains in an observational study of this kind.

Note that the properties in our data do not represent all U.S. rental markets. Results are applicable only to the observed sample.

Findings

1. The effects of offering the Flex service

Multifamily rental properties that gave residents the ability to align their rent payments more closely with the cadence of their income reduced the incidence of recently past-due rents, according to our estimates, thereby reducing resident exposure to late fees and other penalties.

Our findings indicate that making Flex available to residents was associated with longer resident tenure and reduced vacancy rates. We further estimate that, after adjustment, these properties had a higher net operating income relative to the control group.

The clearest evidence is for rent-payment timing: a higher share of current rent and a lower share of rents up to 30 days late. There was no statistically significant effect on rents more than 30 days late, consistent with Flex addressing payment timing rather than fundamental inability to pay.

The resident-stability and operating outcomes point in favorable directions, but they are farther downstream from the Flex mechanism and some of their magnitudes are estimated less precisely.

Housing is the biggest monthly bill for most U.S. households. A program that allows people to tailor the timing of their rent payment to align with their income flow can benefit both renters and properties.

Rethinking rent

Properties offering Flex showed stronger rent reliability, resident stability, and operating performance

	Controlled		Matching	
Outcome	**Estimate**	**95% CI**	**Estimate**	**95% CI**
On-time Payments				
Share of rents current (pp)	+4.5	+1.0 to +8.0	+3.0	+0.1 to +5.9
Share of rents up to 30 days late (pp)	−4.1	−6.6 to −1.6	−2.5	−4.5 to −0.5
Share of rents more than 30 days late (pp)	0.0	−1.6 to +1.6	−0.4	−1.8 to +1.0
Tenure and vacancy				
Median tenure (months)	+6.7	+2.5 to +10.9	+3.7	+0.7 to +6.6
Vacancy rate (pp)	−4.4	−7.7 to −1.1	−2.1	−4.1 to −0.1
Operating performance				
NOI / rent-roll revenue (pp)	+8.5	+1.2 to +15.8	+6.7	−0.2 to +13.6
Collection and write-off costs / rent-roll revenue (pp)	−0.6	−1.0 to −0.2	−0.6	−1.2 to 0.0
Turnover costs / rent-roll revenue (pp)	−3.9	−6.3 to −1.5	−3.2	−5.7 to −0.7
Late fees / rent-roll revenue (pp)	−0.1	−0.3 to +0.1	−0.1	−0.3 to +0.1
Concession gap / rent-roll revenue (pp)	−1.1	−2.1 to −0.1	−0.6	−1.0 to −0.2

Note: Controlled estimates use the property OLS specification; matching estimates use the property matching specification. Confidence intervals are 95 percent intervals. pp = percentage points. "Collection and write-off costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

IMPACT | Improved rent timeliness

AVERAGE EFFECT		
	Controlled estimate (percentage points; 95% CI)	Matched estimate (percentage points; 95% CI)
Share of rents current	+4.5 (+1.0 to +8.0)	+3.0 (+0.1 to +5.9)
Share of rents up to 30 days late	-4.1 (-6.6 to -1.6)	-2.5 (-4.5 to -0.5)
Share of rents more than 30 days late	0.0 (-1.6 to +1.6)	-0.4 (-1.8 to +1.0)

The availability of Flex improved on-time rent payments.

To measure rent-payment timing, we used a sample of 311 properties for which the necessary data were available. We compared properties that offered Flex with properties that did not, using both the controlled specification and a matched comparison. The accompanying table reports 95 percent confidence intervals for both approaches, distinguishing statistical uncertainty from differences between the controlled and matched estimates.

In the controlled estimate, offering Flex increased the share of rents current by 4.5 percentage points. The matched estimate was smaller but directionally consistent, at 3.0 percentage points. Across all properties, the average share of rents up-to-date was 86.7 percent.

This average figure, and the corresponding ones that follow for each reported outcome, provide context for those outcomes; they should not be treated as baselines to which the percentage-point estimates are added, because the estimates measure differences between Flex and otherwise similar non-Flex properties, not changes from the sample average.

Flex availability also reduced short-term delinquency. In the controlled estimate, the share of rents up to 30 days late was 4.1 percentage points lower; in the matched estimate, it was 2.5 percentage points lower. Across all properties, the average share of rents up to 30 days late was 8 percent. We did not detect a statistically significant effect on the share of rents more than 30 days late.

The same pattern appears in a parallel unit-level analysis, reported in full in the appendix. Among 55,977 individual units at 310 properties, Flex availability was associated with higher current-payment rates and lower short-term delinquency.

ANALYSIS: This finding identifies the most direct effect of offering Flex, and is consistent with the basic premise that a flexible rent payment service can help residents stay up to date on their payments.

The results suggest that Flex primarily reduced shorter-term delinquencies by helping residents manage temporary cash-flow disruptions before they developed into more serious payment problems. The

absence of a measurable reduction in rents more than 30 days late is consistent with the idea that flexible payment options can mitigate short-term payment challenges and their downstream consequences, but are less able to address more severe or persistent financial hardship.

The size of the estimated rent-timing effects suggests that enough residents used the service to move property-level payment outcomes. More importantly, the improvement suggests not just access to a new payment option, but real underlying need: Flex appears to address a mismatch between income timing and rent obligations for many households.

IMPACT Extended tenure and vacancy rates

AVERAGE EFFECT		
	Controlled estimate (units shown; 95% CI)	Matched estimate (units shown; 95% CI)
Median tenure (months)	+6.7 (+2.5 to +10.9)	+3.7 (+0.8 to +6.6)
Vacancy rate (percentage points)	-4.4 (-7.7 to -1.1)	-2.1 (-4.1 to -0.1)

The availability of Flex was associated with longer resident tenure and lower vacancy.
To measure resident stability, we examined median resident tenure and vacancy rates at properties offering Flex and at properties that did not. We used both the controlled specification and a matched comparison.

In the controlled estimate, properties offering Flex had median resident tenure that was 6.7 months longer. The matched estimate was smaller but directionally consistent, at 3.7 months. Across all properties in the tenure sample, the average property had a median resident tenure of 24.2 months.

Flex availability was also associated with lower vacancy. In the controlled estimate, vacancy was 4.4 percentage points lower; in the matched estimate, it was 2.1 percentage points lower. Across all properties in the vacancy sample, the average vacancy rate was 12.5 percent, corresponding to an occupancy rate of 87.5 percent.

The same general pattern appears in the unit-level analysis, reported in full in the appendix. Among more than 60,000 individual units in the tenure analysis and more than 72,000 units in the vacancy analysis, Flex availability was associated with longer tenure and lower vacancy.

ANALYSIS: The confidence intervals are consistent with a favorable direction, but the exact magnitudes, especially for vacancy, should be read with more caution than the rent-payment timing results.

These findings are farther downstream from rent-payment timing, so they should be interpreted more cautiously than the rent-timeliness results. Still, they point in a consistent direction: if Flex helps residents

manage temporary payment disruptions, fewer disruptions may translate into longer stays and lower vacancy pressure.

For renters, that can mean greater housing stability and fewer moves following short-term cash-flow shocks. It can be especially important for families with children, helping them remain in familiar schools and avoid disruptions to everyday life. For property operators, longer tenure and lower vacancy can reduce leasing friction, improve revenue predictability, and lower the operating burden associated with replacing residents.

IMPACT	Stronger operating performance

AVERAGE EFFECT		
	Controlled estimate (percentage points; 95% CI)	Matched estimate (percentage points; 95% CI)
NOI / rent-roll revenue	+8.5 (+1.2 to +15.8)	+6.7 (-0.2 to +13.6)
Collection and write-off costs / rent-roll revenue	-0.6 (-1.0 to -0.2)	-0.6 (-1.2 to -0.1)
Turnover costs / rent-roll revenue	-3.9 (-6.3 to -1.5)	-3.2 (-5.7 to -0.7)
Late fees / rent-roll revenue	-0.1 (-0.3 to +0.1)	-0.1 (-0.3 to +0.1)
Concession gap / rent-roll revenue	-1.1 (-2.1 to -0.1)	-0.6 (-1.0 to -0.2)

The availability of Flex was associated with stronger operating outcomes.

To measure property operating performance, we examined net operating income and several related cost categories as shares of rent-roll revenue. Here, too, we used both the controlled specification and a matched comparison.

In the controlled estimate, properties offering Flex had NOI that was 8.5 percentage points higher as a share of rent-roll revenue, though the lower bound of the 95 percent confidence interval implies a much smaller effect of about 1.2 percentage points.[11] The matched estimate was smaller but directionally

[11] The controlled and matching estimates reflect performance aggregated over each property's most recent valid three-month financial reporting period rather than month-by-month variation, and are expressed relative to implied rent-roll revenue over that same interval.

consistent, at 6.7 percentage points. Across all properties in the NOI sample, average NOI was 48.1 percent of rent-roll revenue.[12,13]

Several operating-cost measures moved in the same favorable direction. Turnover costs were 3.9 percentage points lower in the controlled estimate and 3.2 percentage points lower in the matched estimate. Collection-related costs were 0.6 percentage points lower in both approaches. Concessions were also lower, by 1.1 percentage points in the controlled estimate and 0.6 percentage points in the matched estimate. Late-fee revenue was lower as well, though late fees were a small share of rent-roll revenue and should not be interpreted as a major driver of the NOI result.

ANALYSIS: These operating outcomes are farther downstream from rent-payment timing, and the confidence intervals around NOI are wide. The NOI estimates should therefore be read as directional evidence of stronger operating performance, not as precise estimates of expected property-level financial gain.

Reading the evidence as a whole, we think the causal effect of offering Flex on NOI is likely positive, but probably closer to the lower end of the 95 percent confidence interval than to the point estimate. The pre-treatment and placebo evidence supports the direction of the result, but the wide interval, smaller matched estimate, and conservative lower-tail restrictions all counsel caution on magnitude. Those lower-tail restrictions knowingly bias the estimates by removing lower-performing properties; even so, NOI remains favorable in some specifications, though less precise. A milder version of the same caution applies to vacancy, while the rent-timing, tenure, turnover-cost, collection-cost, and concession estimates appear more stable across the evidence.

The cost patterns help explain why NOI may improve when rent-payment timing improves. Fewer short-term payment disruptions can reduce collection activity, lower turnover-related costs, and ease vacancy pressure. Those channels are consistent with the rent-timing and resident-stability findings, but they are not measured with equal precision.

Late-fee revenue also declined, but late fees represent a small share of rent-roll revenue. The late-fee result is therefore better read as evidence of fewer late-payment events than as an important driver of the NOI estimate.

In short, the operating-cost patterns are consistent with Flex creating meaningful savings for operators and stronger outcomes for residents, but the exact size of the NOI effect should be interpreted cautiously.

[12] Net operating income represents the income a property generates after operating expenses but before accounting for taxes, financing costs, and non-operating items such as capital expenditures. As a result, NOI typically is higher than profitability–much higher–yet it remains an informative measure of a property's operational performance and its potential to generate profit.

[13] This figure is roughly consistent with the 54.8 percent NOI as a share of gross potential rent observed across an independent sample of approximately 2,100 properties collected by the National Apartment Association nationally from 2004 to 2021. Shoag, Daniel, and Issi Romem. *Behind the high cost of rent.* Published February 2025. Accessed March 27, 2026. https://www.metrosight.com/articles/behind-the-high-cost-of-rent.

Conclusion

The structure of rent payment in the United States–large, fixed obligations due on a single date each month–creates a harsh mismatch with the timing of household income. For many renters, especially those with limited financial buffers or irregular pay schedules, this misalignment can turn short-term fiscal disruptions into longer-lasting housing instability.

This report shows that addressing that mismatch is associated with meaningful benefits for renters and multifamily operators. Properties that offered Flex had more reliable rent payments, longer resident tenure, and lower vacancy rates, all of which coincided with stronger operating performance. The strongest evidence is for rent-payment timing; the operating outcomes are favorable but their magnitude is estimated less precisely.

The mechanism of flexible rent payments often is overlooked in housing policy discussions. While affordability remains a central challenge, the timing of rent payments relative to income plays its own critical role in undermining housing stability. The study suggests that interventions helping renters align cash flow with financial obligations can also help prevent short-term payment difficulties and consequent housing disruptions, even with no change in the cost of rent.

As flexible rent payment tools become more common, future research using longer property histories and broader geographic coverage could refine these findings. Richer data would make it possible to examine how effects vary by property size, market conditions, resident income, and other property or neighborhood characteristics.

These findings suggest that flexible rent payment tools can be a promising and scalable way to improve rent-payment reliability and support better outcomes for both renters and housing providers. The underlying intervention is straightforward: align rent-payment timing more closely with household cash flow, so temporary timing mismatches are less likely to create payment disruptions, along with their adverse downstream consequences.

Appendix A: Methodology

Data and sample

This study draws on administrative data from Flex, which offers a payment platform that allows residents to split their monthly rent into multiple payments.

These data yield detailed property- and unit-level data on rent collections, financial performance, occupancy, and lease characteristics.

Location and property size. Our sample consists of 488 multifamily rental properties in 25 states. Of these, 134 properties offered the Flex platform as of the date of our analysis (the treated group) and 354 did not (the control group).

Properties range in size from small communities to large complexes, with a median count of 166 units. The sample delivers a broad cross-section of U.S. rental markets, but it is not nationally representative geographically.

The geographic distribution of observed properties

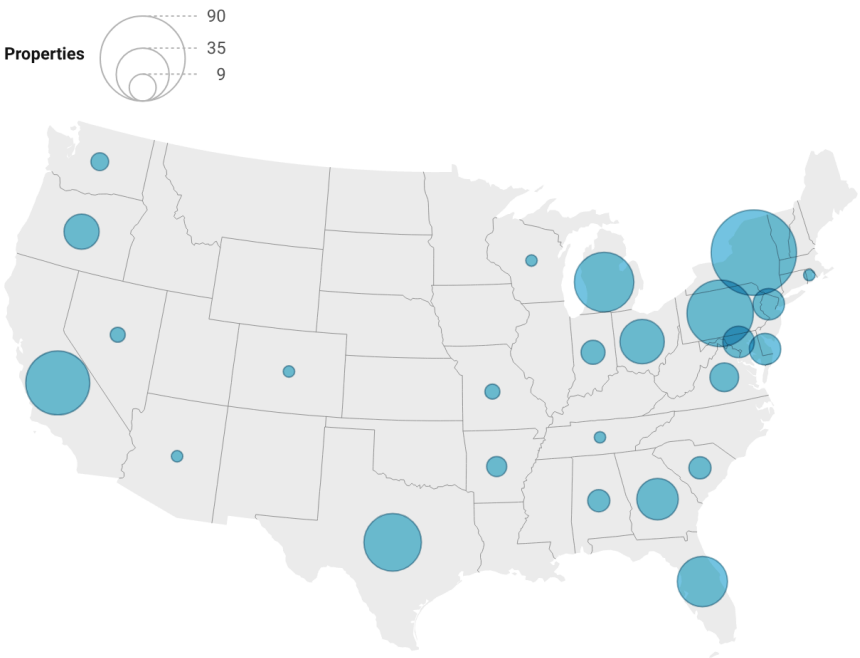

Source: MetroSight

Created with Datawrapper

Property characteristics. In addition to property characteristics such as unit count, average unit size, bedrooms, and rent levels, we supplement property-level data with ZIP-code-level demographic and economic characteristics from American Community Survey 2023 5-year estimates, including median household income, median home value, vacancy rates, unemployment rates, racial composition, and educational attainment. These serve as controls for local market conditions.

Operational data. We construct a cross-sectional dataset using each property's latest available month of data rather than a traditional panel.

This design reflects the structure of the administrative records: Different data sources, such as aged receivables, income statements, occupancy, and rent rolls, report at different frequencies and with different lags. The latest observed month in our data for receivables, occupancy, and concessions was November 2025; for rent rolls it was September 2025.

Because raw data are available for different properties over different time spans, we use each property's most recent observation in each source rather than forcing all properties into a single common month.

This does two things. First, it maximizes sample coverage, because not every property appears in every month of every source. Second, it allows us to measure outcomes at the point when a property has had the longest observed exposure to Flex, which also is when adoption is most likely to have matured. Treatment status is therefore assigned relative to the analysis month for each source.

For example, when we analyze aged-receivables outcomes, we determine whether Flex was live at the property in that property's analysis month—that is, its latest observed month. This property-specific timing rule reduces any mismatch between the timing of the outcome and the timing of Flex availability.

For finance outcomes, we use the property's latest available three-month finance window rather than a single month, and assign treatment relative to that window.

Outcome variables

We examine nine property-level outcomes related to rent collection, financial performance, and occupancy.

Rent collection. We measure the share of units that are current on rent, as well as the share with past-due balances, using aged receivables and rent-roll unit counts. Past-due balances are grouped into minor delinquencies, in which rent is up to 30 days late, and more serious delinquencies, in which rent is more than 30 days late.

Occupancy. We measure resident retention using median tenure in months, constructed from rent-roll and unit data. We measure vacancy as a property-level vacancy rate constructed from the occupancy data, using unit-mix information as the primary source.

All continuous outcomes are winsorized at the 1st and 99th percentiles to limit the influence of extreme values. Because data availability differs across sources, sample sizes vary by outcome family. Rent-payment outcomes are available for 311 properties, finance outcomes for 285, concessions for 140, vacancy for 392, and tenure for 384.

We also estimate unit-level models for whether a unit is current on rent and whether it is vacant. These samples are substantially larger: 55,977 units across 310 properties for rent-payment outcomes and 72,019 units across 391 properties for vacancy, which allows us to control for unit characteristics such as square footage, bedrooms, and rent level in addition to property and market characteristics.

Financial performance. We measure net operating income, turnover costs, collection costs, late fees, and bad debt as shares of implied three-month rent-roll revenue. In the regressions, the financial outcomes' shares subsequently were winsorized at the 1st and 99th percentiles. For each property, we sum the relevant income-statement amounts over a recent three-month finance window and divide by the sum of leased rent for occupied and notice units in the rent roll over the same period.

NOI and late fees are drawn directly from income-statement lines. Turnover costs are constructed from turnover-related ledger entries using the legacy turnover classification logic, which combines account-group mappings with text-based pattern-matching for turnover-related descriptions. Collection costs are defined using account-description rules that capture payment-processing costs, collections-related administrative time, enforcement-specific legal and eviction costs, and realized rent bad-debt losses. Concessions are measured separately from the concessions file as the latest-property-month concession gap, defined as:

$$sum(LEASE\_RENT \ - \ EFFECTIVE\_RENT) \ / \ sum(LEASE\_RENT)$$

Table A1. Outcome summary statistics

Winsorized, sample-specific summary statistics for the main report outcomes.

Outcome	N	Mean	SD	P25	Median	P75
Share Of Rents Up To Date	311	0.867	0.139	0.815	0.915	0.957
Share Late Up To 30 Days	311	0.078	0.079	0.026	0.056	0.106
Share Late More Than 30 Days	311	0.051	0.077	0.004	0.019	0.057
Median Tenure (Months)	384	24.2	17.2	12.9	19.9	31.0
Property Vacancy Rate	392	0.125	0.192	0.030	0.066	0.126
NOI / Rent Roll Revenue	285	0.481	0.275	0.430	0.557	0.629
Turnover / Rent Roll Revenue	285	0.065	0.081	0.021	0.043	0.076
Collection Costs / Rent Roll Revenue	285	0.017	0.021	0.000	0.008	0.029
Late Fees / Rent Roll Revenue	285	0.007	0.008	0.002	0.005	0.010
Concession Gap / Rent Roll Revenue	140	0.019	0.032	0.000	0.000	0.028

Note: The sample excludes properties with missing control variables. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

Treatment definition

We define Flex treatment using a binary property-level availability measure.

A property is treated if the Flex platform was live at the property during the month the property was being analyzed. This definition captures whether Flex was offered to residents at the time of the analysis and is the treatment measure used in the OLS, matching, and unit-level specifications.

Control variables

All specifications include a common set of property- and market-level controls. Property characteristics include total unit count, average leased rent (in logs), average unit square footage (in logs), and average number of bedrooms. Market characteristics, measured at the ZIP code level, include median household income (in logs), median home value (in logs), vacancy rate, unemployment rate, share of residents who are white, and share with a bachelor's degree or higher.

Unit-level specifications additionally control for each unit's rent (in logs), square footage (in logs), bedroom and bathroom counts, and month-to-month lease status, with standard errors clustered at the property level.

All specifications also include state fixed effects to absorb differences in regulatory environments, housing markets, and other factors.

Our reported statistical models do not control for the specific month each property was observed. This decision reflects the nature of the data: rather than tracking the same properties over an extended period, our analysis draws–from a very short history–a cross-sectional snapshot in which each property contributes its most recent month with available data. This approach results in a collection of observation months that are fairly tightly clustered.[14] Given that clustering, adding month controls would do little to control for economic trends or seasonal variation. Instead, it would mainly limit the analysis to be informed only by a comparison of a thinner sample of properties that happen to share the same observation month, a much narrower and less-informative basis.[15] Importantly, our results aren't sensitive to this choice, as described in more detail in the additional robustness checks section below.

[14] In our core models, excluding concessions, more than 99 percent of properties' latest observed months fall within six months of the last month in which each outcome is observed across all properties, and nearly all fall within 12 months.

[15] The informational cost of adding observation-month controls is most evident in the matched-sample design. Once these controls are introduced, the matched samples become smaller and post-match covariate balance worsens. This suggests that the month indicators do less to absorb meaningful time variation than to restrict the comparison to a thinner set of properties observed in the same month.

Table A2. Control summary statistics

Winsorized full-cross-section control summary statistics.

	N	Mean	SD	P25	Median	P75
Number of units	418	178.8	158.9	36.3	165.5	277.8
log (leased rent)	409	7.5	0.6	7.2	7.4	7.8
log (median income)	439	11.3	0.4	11.1	11.3	11.6
log (home value)	438	13.0	0.7	12.4	12.9	13.5
Vacancy rate	439	8.3	5.4	4.5	6.8	10.3
Unemployment rate	439	5.6	2.0	4.0	5.3	7.1
Percent white	439	58.2	21.4	42.0	63.1	75.5
Percent bachelors	439	44.1	19.1	29.3	40.2	60.4
log (average sqft)	400	6.7	0.3	6.6	6.7	6.9
Average bedrooms	415	1.7	0.6	1.4	1.7	2.0

Source: MetroSight • Created with Datawrapper

Empirical strategy and results

We use several complementary observational specifications, including controlled property-level regressions, matching, and unit-level models, along with placebo exercises, sensitivity analyses, and additional robustness checks. These approaches do not eliminate the possibility of unobserved differences between Flex and non-Flex properties, but they help assess whether the same empirical pattern appears after alternative adjustments for observables and in related samples.

The placebo exercises test whether similar patterns appear when the analysis is applied to pre-rollout outcomes, where Flex should not yet have affected the measured results. The sensitivity analyses gauge how large unobserved selection would need to be to explain away the estimated results under stated assumptions, and should not be read as proof that omitted confounding is absent. Finally, we report stress tests that knowingly introduce sample-selection bias in a conservative direction to examine whether the main estimated patterns depend on lower-performing properties being overrepresented among control properties.

Ordinary least squares (OLS)

Our baseline specification regresses each outcome on the binary Flex indicator, the full set of controls, and state fixed effects. We use heteroskedasticity-robust standard errors (HC3), which provide reliable inference in moderately sized samples with potential leverage points. For property outcome y:

$$y_i = \beta \times HasFlex_i + \gamma \times X_i + \delta_{s(i)} + \varepsilon_i$$

Where *y* is the winsorized outcome, *HasFlex* is a treatment indicator, *X* is a vector of property and market controls, and delta represents state fixed effects.

As robustness checks, we estimated models weighted by property size and models restricted to properties with at least 10 units. The unit-count restriction leaves the main results essentially unchanged.

The OLS coefficient on the Flex indicator estimates the average difference in outcomes between treated and control properties, conditional on the included controls. The specification adjusts for observed property, market, and geographic differences, but it does not adjust for unobserved factors such as management quality, operational discipline, resident financial risk, technology orientation, or prior property stabilization. A causal interpretation therefore depends on a selection-on-observables assumption: after conditioning on the included controls, remaining unobserved differences are not driving the outcome gap.

Table A3 shows property-level outcomes of the control approach. In the table, *Beta* denotes the estimated coefficient of interest in the property-level OLS specification, *SE* its standard error, *p* the p-value, and *N* the number of observations.

Table A3. Property OLS

Outcome	Beta	SE	95% CI	p	N
Share Of Rents Up To Date	0.045 **	0.018	0.010 to 0.080	0.013	311
Share Late Up To 30 Days	−0.041***	0.013	-0.066 to -0.016	0.001	311
Share Late More Than 30 Days	0.000	0.008	-0.016 to 0.016	0.984	311
Median Tenure (Months)	6.7***	2.1	2.5 to 10.9	0.002	384
Property Vacancy Rate	−0.044***	0.017	-0.077 to -0.011	0.009	392
NOI / Rent Roll Revenue	0.085**	0.037	0.012 to 0.158	0.024	285
Turnover / Rent Roll Revenue	−0.039***	0.012	-0.063 to -0.015	0.001	285
Collection Costs / Rent Roll Revenue	−0.006***	0.002	-0.010 to -0.002	0.009	285
Late Fees / Rent Roll Revenue	−0.001*	0.001	-0.003 to 0.001	0.099	285
Concession Gap / Rent Roll Revenue	−0.011**	0.005	-0.021 to -0.001	0.022	140

Note: Outcomes are winsorized at 1/99. Specification includes state fixed effects and baseline controls. Standard errors are heteroskedasticity-robust HC3 standard errors. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

The confidence intervals around these estimates should be read alongside the point estimates. They show that the main OLS results generally point in the expected direction, but that the magnitude of some effects is estimated with meaningful uncertainty. This is especially relevant for downstream operating outcomes such as vacancy and NOI.

The vacancy result is directionally favorable, but its confidence interval should be considered when interpreting the practical size of the effect.

For NOI, the point estimate is positive and statistically significant, but its 95 percent confidence interval spans a wide range, with a lower bound implying a much more modest effect than the central estimate.

For that reason, the estimated NOI coefficient should be interpreted as directionally favorable evidence rather than as a precise estimate of expected property-level financial gain.

Inverse probability weighting (IPW-ATT)

To better account for differences between treated and control properties on observed covariates, we estimate the average treatment effect on treated properties (ATT) using inverse probability weighting (IPW).

We first estimate a propensity score, defined as the predicted probability that Flex was live at a property in the relevant source-specific analysis period, using a logistic regression of the treatment indicator on the full set of observed controls and state indicators. Let

$$e_i = P(HasFlex = 1|X, \delta_s)$$

denote this propensity score where again *X* is a vector of covariates and delta state indicators.

For ATT weighting, treated observations receive weight one, while control observations receive weight

$$w_i = \frac{e_i}{1-e_i} .$$

These weights reweight the control group so that, on observed characteristics, it more closely resembles the treated group. This exercise does not control for unobserved differences, and ultimately hinges on the same selection-on-observables assumption as the OLS estimates.

We trim observations with estimated propensity scores below 1 percent or above 99 percent to improve overlap between the treated and control distributions. We also evaluated alternative trimming thresholds of 2 percent, 5 percent, and 10 percent as sensitivity checks, but the primary reported results use IPW-ATT with 1 percent trimming.

After reweighting, we estimate the treatment effect using weighted least squares (WLS):

$$y_i = \beta \times HasFlex_i + \gamma \times X_i + \delta_{s(i)} + \varepsilon_i$$

We assess covariate balance using standardized mean differences (SMDs) before and after reweighting. This balance check is limited to observed covariates; it cannot show that the weighted control group is balanced on unobserved factors. As a rule of thumb, post-weighting SMDs below 0.10 in absolute value often are taken as evidence of good balance on observed covariates, although in our data, balance improves substantially without every covariate necessarily falling below that threshold in every outcome sample.

Table A4 shows property-level outcomes of the matching approach. In the table, *ATT* denotes the average treatment effect on the treated, which corresponds to the Beta coefficient estimate from the property-level WLS specification. *SE* denotes that coefficient's standard error, *p* the p-value, *N kept* the number of observations retained after trimming, *Treated* and *Control* the group sample sizes, and *Post |SMD|* the post-matching standardized mean difference.

Table A4. Property matching

Outcome	ATT	SE	95% CI	p	N Treat/Ctrl	Post |SMD|
Share Of Rents Up To Date	0.030 **	0.015	0.001 to 0.059	0.048	129/137	0.089
Share Late Up To 30 days	−0.025***	0.010	-0.045 to -0.005	0.010	129/137	0.089
Share Late More Than 30 Days	−0.004	0.007	-0.018 to 0.010	0.543	129/137	0.089
Median Tenure (Months)	3.7**	1.5	0.8 to 6.6	0.015	126/171	0.069
Property Vacancy Rate	−0.021**	0.010	-0.041 to -0.001	0.037	125/176	0.063
NOI / Rent Roll Revenue	0.067*	0.035	-0.002 to 0.136	0.057	128/119	0.048
Turnover / Rent Roll Revenue	−0.032**	0.013	-0.057 to -0.007	0.012	128/119	0.048
Collection Costs / Rent Roll Revenue	−0.006**	0.003	-0.012 to -0.001	0.027	128/119	0.048
Late Fees / Rent Roll Revenue	−0.001	0.001	-0.003 to 0.001	0.402	128/119	0.048
Concession Gap / Rent Roll Revenue	−0.006***	0.002	-0.010 to -0.002	0.003	56/80	0.084

Note: Matching results use the IPW-ATT specification with a 1% trim. Post |SMD| is mean absolute SMD after weighting/matching. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

The confidence intervals should again be read alongside the point estimates. The matching estimates generally point in the same direction as the OLS estimates, but some are estimated with meaningful uncertainty.

This is especially important for NOI: the point estimate is positive, but it is only statistically distinguishable from zero at the 10 percent level, and its 95 percent confidence interval includes no effect. The matching NOI result should therefore be interpreted as directionally favorable supporting evidence, not as a precise estimate of expected property-level financial gain.

Unit-level analysis

We complement the property-level analysis with unit-level regressions for four outcomes: whether a unit is current on rent, whether it is vacant, unit tenure in months, and the unit-level concession gap.

These specifications mirror the property-level OLS and IPW-ATT models, but add unit-level controls for log rent, log square footage, the number of bedrooms and bathrooms, and month-to-month lease status, alongside property and market controls. Standard errors are clustered at the property level to account for within-property correlation across units.

$$y_i = \beta \times HasFlex_i + \lambda \times Z_i + \gamma \times X_i + \delta_{s(i)} + \varepsilon_i$$

where *y* still represents the outcome of interest, *X* a vector of property level controls, delta the state fixed effects, and now *Z*, a vector of unit-level controls.

The unit-level analysis serves three purposes. First, it provides a robustness check at a finer observational level. The reported samples include 55,977 units across 310 properties for the rent-current outcome, 72,019 units across 391 properties for vacancy, 60,299 units across 384 properties for tenure, and 27,290 units across 138 properties for concessions.

Second, it allows us to account for unit-level heterogeneity, such as differences in rent, size, bedrooms, bathrooms, and lease status, that is averaged out in the property-level analysis.

Third, for rent-current outcomes, it provides a mechanism check: if Flex improves rent-payment timing at the property level, we should also see the pattern reflected in unit-level payment status.

Because Flex availability is still defined at the property level, these specifications do not resolve property-level selection; unobserved property or operator differences may still contribute to the estimated unit-level gaps.

Tables A5 and A6 show outcomes based on unit-level analysis for the control and matching approaches. Here, *Beta* denotes the coefficient estimate from the unit-level OLS and WLS specifications.

Table A5. Unit OLS

Outcome	Beta	SE	95% CI	p	N units/props
Unit Up To Date	0.048**	0.023	0.003 to 0.093	0.038	55,977/310
Unit Late Up To 30 Days	−0.051**	0.021	-0.092 to -0.010	0.015	55,977/310
Unit Late More Than 30 Days	0.003	0.009	-0.015 to 0.021	0.744	55,977/310
Unit Vacancy Flag	−0.040**	0.019	-0.077 to -0.003	0.036	72,019/391
Concession Gap (%)	−0.010***	0.003	-0.016 to -0.004	<0.001	27,290/138
Unit Tenure (Months)	7.0***	2.4	2.3 to 11.7	0.004	60,299/384

Note: Baseline unit controls include rent/sqft/bed-bath and property/ZIP controls; state fixed effects included. Standard errors are clustered at the property level. N units/props reports unit count followed by property count. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively.

Source: MetroSight • Created with Datawrapper

Table A6. Unit matching

Outcome	ATT	SE	95% CI	p	N Treat/Ctrl	Post |SMD|
Unit Up To Date	0.029 *	0.017	-0.004 to 0.062	0.097	23,163/22,381	0.110
Unit Late Up To 30 Days	−0.030**	0.014	-0.057 to -0.003	0.029	23,163/22,381	0.110
Unit Late More Than 30 Days	0.001	0.011	-0.021 to 0.023	0.948	23,163/22,381	0.110
Unit Vacancy Flag	−0.018*	0.010	-0.038 to 0.002	0.063	26,971/31,645	0.062
Concession Gap (%)	−0.006***	0.002	-0.010 to -0.002	<0.001	11,238/12,183	0.099
Unit Tenure (Months)	5.6***	1.8	2.1 to 9.1	0.002	24,752/25,250	0.070

Note: Matching results use the same method family and 1% trim as the property baseline (IPW-ATT, trim=0.01). Post |SMD| is mean absolute SMD after weighting/matching. N Treat/Ctrl reports treated unit count followed by control unit count. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively.

Source: MetroSight • Created with Datawrapper

Interpreting results across methods

The applied estimation strategies are necessarily observational and rest on selection-on-observables assumptions. The available data do not support a credible within-property panel design, so the control, matching, and unit-level approaches instead use different cross-sectional comparison structures. All three depend on observed controls and state fixed effects rather than randomized assignment. They therefore improve comparability, but do not by themselves rule out unobserved selection.

No single method is definitive. We interpret the results as a body of evidence: when the control, matching, and unit-level approaches point in similar directions, that convergence does not prove causality, but it increases confidence in the empirical pattern.

Identification challenges and robustness checks

One limitation of this study is that properties offering Flex were not randomly selected. Two selection concerns are especially relevant.

- **Potential understatement bias due to resident stress:** Flex may have been offered at properties where residents were more likely to face financial strain, and therefore had greater need for the service. Those properties may have differed from non-Flex properties in underlying delinquency risk, resident stability, collection challenges, or other characteristics associated both with the decision to offer Flex and with subsequent outcomes. If so, the estimates could understate Flex's effect, because the treated properties would have started from a more difficult resident-payment environment.

- **Potential overstatement bias due to operational strength:** Selection could also run in the opposite direction. Flex properties may have been on stronger operational trajectories, or they may have had management teams with greater operational capacity, technology orientation, or willingness to adopt new resident-facing services. Those same characteristics could also be associated with stronger subsequent operating outcomes independent of Flex. If so, the estimates could overstate Flex's effect, because some of the treated-control difference would reflect pre-existing operational advantages rather than Flex itself.

Because the analysis relies primarily on a cross-sectional snapshot rather than a long panel observed before and after rollout, these concerns cannot be addressed through a conventional event-study or difference-in-differences framework.

To mitigate these risks, the analysis combines several complementary approaches. First, the core specifications control for a rich set of observable property and market characteristics. Second, we apply matching and reweighting methods designed to improve comparability between Flex and non-Flex properties. Third, we conduct pre-treatment balancing checks, placebo and sensitivity exercises that test whether the results appear consistent with simple baseline selection effects and how sensitive they are to omitted-variable concerns. No single exercise is definitive on its own. Rather, agreement across these approaches is best read as supportive convergence across imperfect observational checks, not as independent causal validation or as proof that unobserved selection has been fully eliminated.

Pre-treatment outcome differences

One way to assess the possibility of omitted-variable bias is to compare properties on the main outcome variables before Flex rollout. If properties that later adopted Flex already differed systematically from the control group prior to adoption, those differences could contribute to the estimated treatment effects observed later in the sample.

In this setting, however, conducting a clean pre-treatment comparison is difficult for several reasons. First, we often do not observe Flex properties far enough before rollout to construct a long and uniform pre-treatment history. In many cases, the earliest observed month is at or close to the start of treatment, so the first observed value can serve only as an imperfect proxy for a truly pre-treatment value.

Second, properties that never offered Flex do not have a natural rollout date, so there is no obvious "before Flex" period for them. The control-group baseline must therefore be constructed from a chosen comparison window. That timing choice is central: if the control window is too broad, or mismatched with the treated properties' pre-Flex window, then seasonal trends, different market conditions, or early lease-up and pre-stabilization stages could make the control baseline non-comparable to the treated baseline.

To address these limitations, we roll both groups back to a shared calendar window with substantial overlap, December 2024 through February 2025, and compare outcomes for the reduced sample observed within that common period. For properties that later adopted Flex, we use the most recent strictly pre-rollout observation available within the window. For the control group, we similarly restrict observations to that same period.

Unfortunately, this careful timing rule also makes the sample much smaller than the main cross-sectional sample. The resulting pre-treatment comparison should therefore be read as a limited baseline check. It is informative about baseline comparability among properties for which a reasonably timing-comparable comparison can be constructed, but the feasible pre-treatment sample is too small and potentially selected to support panel approaches or to validate or invalidate the full treatment-period analysis by itself.

Tables A7 and A8 show the results of the exercise.

Across both the control and matched samples, the estimated pre-treatment differences are mostly statistically indistinguishable from zero at conventional levels, even though the main regression framework later detects sizeable and statistically significant differences after Flex rollout. Vacancy rates and late-rent shares tended to be somewhat higher among non-Flex properties, while in contrast some measures of financial performance appeared somewhat stronger at non-Flex properties. In several cases, these differences approach conventional significance thresholds, but the overall pattern is considerably weaker and less consistent than the treatment-period results reported in the main analysis. Even so, we cannot fully rule out the possibility that pre-treatment differences contributed to part of the estimated effects.

Table A7. Property OLS pre-treatment outcome balancing

Outcome	Treated mean	Control mean	Difference	SE	p	N Treat/Ctrl
Share Of Rents Up To Date	0.939	0.923	0.015	0.012	0.214	60/36
Share Late Up To 30 Days	0.048	0.060	−0.012	0.009	0.209	60/36
Share Late More Than 30 Days	0.013	0.017	−0.003	0.004	0.434	60/36
Median Tenure (Months)	31.9	30.3	1.6	2.8	0.564	60/38
Property Vacancy Rate	0.044	0.085	−0.041	0.027	0.131	60/39
NOI / Rent Roll Revenue	0.571	0.600	−0.029	0.044	0.509	58/32
Turnover / Rent Roll Revenue	0.041	0.040	0.002	0.013	0.900	58/32
Collection Costs / Rent Roll Revenue	0.001	0.001	0.000	0.001	0.834	58/32
Late Fees / Rent Roll Revenue	0.003	0.005	−0.002*	0.001	0.067	58/32
Concession Gap / Rent Roll Revenue	0.000	0.000	0.000	0.000	0.163	49/30

Note: Balancing estimates are based on pre-rollout outcomes for Flex properties, observed at each property's most recent pre-rollout month within the December 2024 to February 2025 window. Control properties are observed in the earliest available month within the same window. N Treat/Ctrl reports treated property count followed by control property count. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

Table A8. Property matching pre-treatment outcome balancing

Outcome	Treated mean	Control mean	Difference	SE	p	N Treat/Ctrl
Share Of Rents Up To Date	0.942	0.923	0.019	0.012	0.117	54/36
Share Late Up To 30 Days	0.046	0.060	−0.014	0.009	0.132	54/36
Share Late More Than 30 Days	0.012	0.017	−0.005	0.004	0.279	54/36
Median Tenure (Months)	32.3	30.7	1.7	2.8	0.547	58/37
Property Vacancy Rate	0.043	0.068	−0.025*	0.014	0.095	58/38
NOI / Rent Roll Revenue	0.570	0.603	−0.034	0.044	0.446	56/31
Turnover / Rent Roll Revenue	0.040	0.038	0.002	0.013	0.875	56/31
Collection Costs / Rent Roll Revenue	0.001	0.001	0.000	0.001	0.859	56/31
Late Fees / Rent Roll Revenue	0.003	0.005	−0.002*	0.001	0.096	56/31
Concession Gap / Rent Roll Revenue	0.000	0.000	0.000	0.000	0.105	48/30

Note: Balancing estimates are based on pre-rollout outcomes for Flex properties, observed at each property's most recent pre-rollout month within the December 2024 to February 2025 window. Control properties are observed in the earliest available month within the same window. N Treat/Ctrl reports treated property count followed by control property count. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

This pattern also does not rule out the possibility that unobserved differences between properties contributed to the estimated effects. In particular, properties adopting Flex may have differed from the control group in ways that are difficult to observe directly. Bias could run in either direction: greater resident financial stress would tend to make Flex properties look worse at baseline, while stronger operations would tend to make them look better. The pre-treatment comparisons for vacancy make the operational-strength concern worth considering, but the baseline evidence is limited, largely statistically insignificant, not uniform across outcomes, and drawn from a smaller timing-comparable sample. These factors could have influenced both the decision to adopt Flex and subsequent operating outcomes independent of the service itself.

At the same time, the absence of large and consistently significant pre-treatment differences across the primary outcomes provides additional evidence against the view that the headline findings are driven entirely by simple baseline selection on observed outcome levels. A related placebo exercise using the full regression framework is described later in this appendix.

Placebo robustness check

As a placebo exercise, we repeat the property-level OLS and matching analyses using the same timing-comparable pre-treatment samples described above. Intuitively, pre-rollout outcomes should not be affected by the subsequent introduction of Flex. If, on the other hand, the estimated effects in the main specification were instead driven by pre-existing differences between Flex and control properties, we would expect to observe similar patterns in these placebo regressions.

Unlike the balancing-table comparisons above, the placebo regressions retain the full set of controls and fixed effects from the main specification.

As noted, pre-treatment data are not available for all Flex properties. Among those with such data, all but three observations fall within the period from December 2024 through February 2025. To avoid introducing bias from differential timing, and following the same sample construction logic as in the pre-treatment balancing exercise, we restrict the control group to this same window and exclude both Flex and control properties that cannot be mapped into it. We then estimate the same specification as in the baseline model, including the full set of controls, state fixed effects, and winsorization.

The OLS placebo estimates are generally small in magnitude and statistically indistinguishable from zero. In several cases, including net operating income, the estimated coefficients take the opposite sign from those in the main specification. The only outcome that remains statistically significant at the 5 percent level is late fees, which represent only a minimal share of the overall effect on operating gains associated with Flex. Taken together, these results are consistent with the view that the main findings are not simply explained by observed pre-treatment differences.

We also conduct a parallel placebo exercise using the matching specification applied in the main analysis, maintaining the same one percent trimming threshold, control variables, state handling, and one-row-per-property cross-sectional structure. In this case, future-Flex properties are represented by a single placebo observation equal to their most recent strictly pre-rollout value within the December 2024 to February 2025 window. Control properties are similarly restricted and represented by a single observation per property, selecting December 2024 where available, otherwise January 2025, and otherwise February 2025. Under this restriction, the placebo samples range from 78 to 96 properties depending on the outcome.

Table A9. Property OLS placebo

Outcome	Beta	SE	p	N
Share Up To Date	0.017	0.012	0.173	96
Share Late Up To 30 Days	−0.016*	0.009	0.087	96
Share Late More Than 30 Days	0.000	0.005	0.971	96
Median Tenure (Months)	1.1	3.1	0.724	98
Property Vacancy Rate	−0.015	0.012	0.214	99
NOI / Rent Roll Revenue	−0.010	0.033	0.768	90
Turnover / Rent Roll Revenue	0.006	0.004	0.164	90
Collection Costs / Rent Roll Revenue	0.000	0.000	0.701	90
Late Fees / Rent Roll Revenue	−0.001**	0.001	0.032	90
Concession Gap / Rent Roll Revenue	0.000	0.000	0.768	79

Note: Placebo estimates are based on pre-rollout outcomes for Flex properties, observed at each property's most recent pre-rollout month within the December 2024 to February 2025 window. Control properties are observed in the earliest available month within the same window. Outcomes are winsorized at 1/99. Specification includes state fixed effects and baseline controls. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

Table A10. Property matching placebo

Outcome	ATT	SE	p	N Treat/Ctrl	Post |SMD|
Share Of Rents Up To Date	0.011	0.014	0.441	54/36	0.143
Share Late Up To 30 Days	−0.014	0.009	0.118	54/36	0.143
Share Late More Than 30 Days	0.004	0.008	0.654	54/36	0.143
Median Tenure (Months)	−0.3	4.0	0.946	58/37	0.155
Property Vacancy Rate	−0.007	0.011	0.506	58/38	0.155
NOI / Rent Roll Revenue	−0.007	0.028	0.802	56/31	0.084
Turnover / Rent Roll Revenue	0.006*	0.003	0.058	56/31	0.084
Collection Costs / Rent Roll Revenue	0.000	0.000	0.365	56/31	0.084
Late Fees / Rent Roll Revenue	−0.001	0.001	0.306	56/31	0.084
Concession Gap / Rent Roll Revenue	0.000	0.000	0.599	48/30	0.105

Note: Placebo estimates are based on pre-rollout outcomes for Flex properties, observed at each property's most recent pre-rollout month within the December 2024 to February 2025 window. Control properties are observed in the earliest available month within the same window. Matching results use the IPW-ATT specification with a 1% trim. Post |SMD| is mean absolute SMD after weighting/matching. N Treat/Ctrl reports treated property count followed by control property count. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. "Collection costs" refers to rent collection and write-off costs.

Source: MetroSight • Created with Datawrapper

Relative to the main matching results, the placebo estimates are modest in magnitude, and all but one are statistically indistinguishable from zero in the direction of the main effects.[16] Because these samples are smaller and inherit the same timing-comparability restrictions as the pre-treatment balancing exercise, they should be read as limited-sample placebo checks. They provide additional evidence that the main estimates are not obviously mirrored by observed pre-treatment differences in that subset, but they do not fully rule out selection concerns in the treatment-period estimates.

Tables A9 and A10 show property-level outcomes of the placebo exercises for the control approach and the matching approach.

Conservative lower-tail sample restriction

This check returns to the second selection concern described above: potential overstatement bias due to operational strength. If Flex properties were already stronger operators, or if non-Flex control properties included a disproportionate number of low-performing properties, then part of the estimated Flex-control difference could reflect pre-existing operational weakness in the control group rather than Flex itself.

Unlike the pre-treatment balance and placebo sections, this exercise uses the treatment-period sample. We re-estimate the main property-level OLS and matching specifications after excluding properties below selected occupancy and NOI thresholds observed in that sample.

These restrictions are not preferred specifications because they condition on outcomes that are either dependent variables themselves or closely related to them, and therefore knowingly introduce sample-selection bias. In the screens considered here, the bias runs in a conservative direction: because the main estimates imply stronger outcomes for Flex properties, removing low-occupancy or low-NOI properties, which are relatively more prevalent in the control group, tends to make the control group look better and therefore tends to understate the Flex-control difference. This bias is most direct for vacancy and NOI: the screens condition on occupancy, which is the inverse of vacancy, and on NOI-related performance itself. For rent timing, tenure, turnover costs, collection costs, and concessions, the restrictions are still relevant but less directly tied to the dependent variable. The more restrictive screen should introduce more of this bias.

Table A11 reports two restricted versions of the analysis. The first excludes properties with negative NOI and occupancy below 50 percent. The second applies a more restrictive screen, excluding properties with NOI below 50 percent of rent-roll revenue and occupancy below 90 percent.

Because these restrictions both reduce sample size and intentionally bias the estimates against the main treatment-period differences, weaker estimates in Table A11 should be expected. Even so, the main rent-payment timing results remain directionally favorable under both restrictions. The share of rents up to date remains positive, and the share of rents up to 30 days late remains negative. Under the stricter restriction, both OLS rent-payment estimates remain statistically significant, while the matching estimates remain directionally consistent but are less precise for current rent. Median tenure and turnover costs also remain directionally consistent with the main results, with stronger persistence in OLS and somewhat weaker but still generally supportive matching estimates.

[16] The estimate with respect to turnover cost as a share of rent roll revenue was statistically significant at the 10 percent level, but in the opposite direction of the main results.

Table A11. Lower-tail sample restriction checks

Outcome	Model	Canonical estimate (N)	Exclude negative NOI & occupancy <50% (N)	Exclude NOI/RR <50% & occucpancy <90% (N)
Share Of Rents Up To Date	OLS	+0.045** (311)	+0.041* (269)	+0.037* (239)
	Matching	+0.030** (266)	+0.035* (230)	+0.024 (203)
Share Late Up To 30 Days	OLS	−0.041*** (311)	−0.027*** (269)	−0.029*** (239)
	Matching	−0.025*** (266)	−0.020* (230)	−0.019** (203)
Median Tenure (Months)	OLS	+6.671*** (384)	+6.097*** (278)	+4.971** (248)
	Matching	+3.748** (299)	+3.985** (242)	+2.876* (214)
Property Vacancy Rate	OLS	−0.044*** (392)	−0.026** (278)	−0.012* (248)
	Matching	−0.021** (302)	−0.014* (242)	−0.007 (214)
NOI / Rent Roll Revenue	OLS	+0.085** (285)	+0.049 (278)	+0.007 (248)
	Matching	+0.068** (248)	+0.050* (238)	+0.027 (212)
Turnover / Rent Roll Revenue	OLS	−0.039*** (285)	−0.033*** (278)	−0.022*** (248)
	Matching	−0.033*** (248)	−0.029** (238)	−0.019** (212)

Note: Parentheses in column headers indicate that parenthetical values report sample sizes. One, two, or three asterisks denote 10%, 5%, and 1% statistical significance, respectively. RR = rent-roll revenue.

Source: MetroSight • Created with Datawrapper

The downstream operating outcomes most directly implicated by the restrictions are more sensitive, as expected. Vacancy remains negative and statistically significant in OLS even under the stricter restriction, but the matching vacancy estimate weakens and is no longer statistically distinguishable from zero. NOI weakens more materially: the less restrictive matching estimate remains positive and statistically significant, but the OLS estimate and the stricter matching estimate are not statistically distinguishable from zero.

Null or weaker estimates in this exercise should not be read as evidence that the corresponding treatment-period effects are zero; the restricted samples are smaller and intentionally biased against the main treatment-period differences. Rather, the exercise shows that the core rent-timing, tenure, and turnover patterns persist even after introducing conservative sample-selection bias, while NOI and, to a lesser extent, vacancy remain more sensitive and should be interpreted cautiously, especially as estimates of magnitude.

Sensitivity to unobserved confounding

A central concern in any observational study is the potential for unobserved factors to confound the relationship between treatment and outcomes. We assess the robustness of our estimates to this concern using the Oster sensitivity framework, which builds on the earlier intuition of Altonji et al.[17,18]

The idea is straightforward. We compare two regressions: a raw regression of the outcome on the treatment indicator alone, with no controls, and a full regression that includes all controls and state fixed effects. The change in the treatment coefficient between the two specifications reflects the degree to

[17] Oster E. Unobservable selection and coefficient stability: theory and evidence. *J Bus Econ Stat.* 2019;37(2):187–204. https://www.tandfonline.com/doi/full/10.1080/07350015.2016.1227711
[18] Altonji JG, Elder TE, Taber CR. Selection on observed and unobserved variables: assessing the effectiveness of Catholic schools. *J Polit Econ.* 2005;113(1):151–184. https://www.nber.org/papers/w7831

which observed controls absorb selection. If adding the full set of controls barely moves the coefficient, this suggests that unobserved factors, if they are related to the outcome in similar ways as the observed controls, would also have limited impact. The exercise should be read as a structured sensitivity benchmark, not as a test that the relevant omitted factors are absent.

The Altonji ratio formalizes this intuition by comparing how much the estimated coefficient changes when controls are added. A ratio of 3, for example, suggests that selection on unobservables would need to be three times as strong as selection on all observed variables combined to fully explain away the estimated effect.

The Oster delta extends this idea by accounting for changes in R-squared. Assuming that the maximum attainable R-squared–if all relevant variables were observed–is 1.3 times the R-squared of the full model, the Oster delta reports the minimum strength of selection on unobservables relative to observables needed to drive the coefficient to zero. A delta greater than one in absolute value indicates that unobserved confounders would need to be more important than the full set of observed controls to eliminate the estimated effect.

In our application, much of the observed-selection benchmark's explanatory-power gain comes from state fixed effects, so the exercise should be read as benchmarking unobserved selection against the full observed adjustment set, not only against measured property and market covariates.

In our results, the reported property-level outcomes mostly pass this benchmark. Under the main Oster calibration, the rent-payment timing results are especially strong: the rent-current and short-delinquency estimates would require unobserved selection roughly 58 and 32 times as strong as the full observed-selection benchmark, respectively, to be explained away. Several operating outcomes also show meaningful robustness: turnover costs, NOI, median tenure, vacancy, and collection costs would require unobserved selection of roughly 19, 8, 7, 5, and 3 times that benchmark, respectively.

The remaining reported outcomes warrant more caution. Late fees pass both calibrations but with little headroom, while concessions pass the main calibration only narrowly and do not pass the stricter calibration that sets the maximum R-squared to 1.0 (as opposed to just 1.3 times the R-squared of the full model). The exercise is also not supportive for rents more than 30 days late, where the main analysis itself does not estimate a meaningful effect. This pattern is consistent with the rest of the evidence: the strongest support is for rent-payment timing, tenure, and related operating outcomes, while the most severe delinquency margin and some secondary outcomes should be interpreted more cautiously.

Additional robustness checks

We subjected the main results to a broad set of robustness checks that vary the sample, weighting, control set, matching design, and inference method. On the OLS side, we re-estimated the models after restricting the sample to properties with at least 10 units and after weighting observations by property size so that larger assets receive proportionally greater influence. The minimum-unit restriction leaves the headline results essentially unchanged. Unit weighting makes some estimates less precise, especially for vacancy and, to a lesser extent, rent timeliness and NOI, but it does not reverse the overall pattern: properties offering Flex continue to show better rent timeliness, longer resident tenure, lower turnover costs, lower collection-related costs, lower concessions, and higher NOI. We also ran a drop-one-control exercise across the seven core property outcomes; no single covariate drives the findings, as the coefficients retain the same sign throughout. Finally, excluding fully vacant properties leaves the vacancy result essentially unchanged.

We also tested whether the results depend on a particular matching specification. The matching analysis was re-estimated using both inverse-probability-weighted ATT and doubly robust AIPW ATT estimators, with trimming thresholds of 1, 2, 5, and 10 percent. Across these alternatives, the core outcomes overwhelmingly preserve the expected sign, and most remain statistically distinguishable from zero at conventional or near-conventional levels. Covariate balance improves materially after weighting, with standardized mean differences falling well below their raw treated-control gaps, indicating substantially better overlap on observed covariates between treated and control properties after reweighting.

As an additional inference check, we also re-estimated the property-level OLS models with standard errors clustered at the state level rather than using only heteroskedasticity-robust HC3 errors. Because the sample includes 25 states, we view this as a conservative sensitivity check rather than the primary basis for inference. Even under that stricter treatment of standard errors, the main rent-timeliness, tenure, NOI, turnover, and collection-cost results remain statistically meaningful and retain the same sign; vacancy and concessions become less precisely estimated, but they do not reverse. Taken together with the unit-level models and the Altonji-Oster sensitivity calculations, these checks all point in the same direction: not every secondary estimate is equally precise in every specification, but the central empirical pattern is durable.

As noted earlier, our results aren't sensitive to the inclusion of month-of-observation controls. All core estimates maintain the same direction, and some, particularly net operating income, become larger. The substantive conclusions are unchanged. Restricting the control sample to observations within 12 months of the latest month does not change the direction or statistical significance of the core results. Even under a more restrictive six-month observation period, all core results hold with the exception of the already-weak late-fee revenue result, which loses significance at the 10 percent level.

Disclosure

The authors would like to thank and acknowledge Flexible Finance, Inc. for funding this research and providing access to the data used.

About

MetroSight is a San Francisco Bay Area-based economic research firm specializing in data-driven analysis of housing, labor, regional and urban economics.