

Time to Recidivism After Release

N = 432 · Dataset: recidivism.csv

Statistical Tests Used

Objective	Test	Variables	Why this test
To describe the weeks_to_arrest and arrested indicator among the released prisoners.	Descriptive Statistics	weeks_to_arrest	Summary statistics for the continuous variables.
To describe the weeks_to_arrest and arrested indicator among the released prisoners.	Frequency & Percentage Distribution	arrested	Frequency and percentage distribution for the categorical variables.
To estimate survival curves for the financial_aid vs no-aid groups (Kaplan-Meier).	Kaplan-Meier Survival Estimate	weeks_to_arrest, arrested, financial_aid	Descriptive survival curves stratified by 'financial_aid' — median survival and event/censoring counts.
To estimate survival curves for the financial_aid vs no-aid groups (Kaplan-Meier).	Log-Rank Test	weeks_to_arrest, arrested, financial_aid	Nonparametric comparison of survival across levels of 'financial_aid'.
To estimate survival curves for the financial_aid vs no-aid groups (Kaplan-Meier).	Cox Proportional Hazards Regression	weeks_to_arrest, arrested, financial_aid	Hazard-ratio estimates for 1 predictor(s), with the proportional-hazards assumption checked.
To model how age and prior offenses predict recidivism time (Cox regression).	Kaplan-Meier Survival Estimate	weeks_to_arrest, arrested	Descriptive survival curves — median survival and event/censoring counts.
To model how age and prior offenses predict recidivism time (Cox regression).	Cox Proportional Hazards Regression	weeks_to_arrest, arrested, age, prio	Hazard-ratio estimates for 2 predictor(s), with the proportional-hazards assumption checked.


Methods

All statistical analyses were performed in Python using the SciPy and statsmodels libraries via thericerca's deterministic analysis engine (v0.1.0). Statistical significance was evaluated at $\alpha = .05$ (a 95% confidence level). Descriptive statistics were computed to summarise the study's variables. For each inferential test, the appropriate effect size is reported with its 95% confidence interval, alongside the achieved statistical power to detect the observed effect (Cohen, 1988).

Results

To describe the weeks_to_arrest and arrested indicator among the released prisoners.

This simply summarises the data — the average, spread, and range of each variable are in the table above.




Distribution of weeks_to_arrest.

max	mean	median	min	mode	n	range	std	variable	variance
52	45.854	52	1	52	432	51	12.662	weeks_to_arrest	160.334

To describe the weeks_to_arrest and arrested indicator among the released prisoners.

This shows how the responses are distributed: how many people fall into each category, and what share of the total that is.




Frequency of each arrested category.

count	percent	value	variable
318	73.611	0	arrested
114	26.389	1	arrested

To estimate survival curves for the financial_aid vs no-aid groups (Kaplan-Meier).

We estimated how weeks to arrest plays out over the follow-up window — for each moment, the share of people who haven't yet experienced the event. The table above gives median times and event counts.

Statistically: KM: N = 432, 114 event(s) across 2 group(s); see per-group medians in the table



Kaplan-Meier survival curves by financial_aid


censored	events	group	median_survival	n	restricted_mean
150	66	0	—	216	0.859

censored	events	group	median_survival	n	restricted_mean
168	48	1	—	216	0.889

To estimate survival curves for the financial_aid vs no-aid groups (Kaplan-Meier).

We compared how quickly the event happens across the financial aid groups. They look similar in this sample.

Statistically: Log-rank $\chi^2(1, N = 432) = 3.84, p = .050$



Kaplan-Meier survival curves by financial_aid

events	group	n
66	0	216
48	1	216

To estimate survival curves for the financial_aid vs no-aid groups (Kaplan-Meier).

We modelled how financial aid shape the hazard of the event over time. In this sample, they don't clearly change when the event occurs.


Statistically: Cox regression: $N = 432$ (114 event(s)), concordance = .55, LR test $p = .050$

coef	hazard_ratio	hr_ci_high	hr_ci_low	p_value	predictor
-0.369	0.691	1.003	0.477	0.052	financial_aid

To model how age and prior offenses predict recidivism time (Cox regression).

We estimated how weeks to arrest plays out over the follow-up window — for each moment, the share of people who haven't yet experienced the event. The table above gives median times and event counts.

Statistically: KM: N = 432, 114 event(s), median survival = not reached



Kaplan-Meier survival curve

censored	events	group	median_survival	n	restricted_mean
318	114	overall	—	432	0.884

To model how age and prior offenses predict recidivism time (Cox regression).

We modelled how age and prio shape the hazard — the moment-to-moment risk of the event. The predictors do carry information about when the event happens. Each predictor's hazard ratio in the table shows the direction and size of its effect.

Statistically: Cox regression: N = 432 (114 event(s)), concordance = .63, LR test $p < .001$

coef	hazard_ratio	hr_ci_high	hr_ci_low	p_value	predictor
-0.069	0.933	0.972	0.896	0.001	age
0.095	1.099	1.159	1.042	0	prio

⚠ Proportional-hazards assumption may be violated; hazard ratios should be interpreted cautiously.

Discussion

How weeks to arrest differs across groups of financial aid. The Kaplan-Meier estimate describes how weeks to arrest unfolds across the sample: at each observed moment, the survival curve shows the share of people who have not yet experienced the event. It handles right-censoring correctly (people whose event hasn't happened by the end of follow-up are treated as still at risk, not as non-events). Read the median and per-group event counts in the table as the anchor points; the curve itself reveals the **shape** of the timing (early drops vs. late attrition) that a single number cannot convey.

Whether weeks to arrest differs across groups of financial aid. The log-rank test did not find a statistically reliable difference in weeks to arrest across the financial aid groups. The survival curves may still cross or drift apart at different points — worth reading the plot alongside — but the overall pattern is consistent with the groups sharing the same timing.

How financial aid predict weeks to arrest. The Cox model with financial aid did not reliably predict when the event happens; the overall likelihood-ratio test was not significant. Individual predictors may still be worth inspecting in the hazard-ratio table, but the combined story is consistent with these variables not shaping the event timing in this sample.

The distribution of weeks to arrest until the event. The Kaplan-Meier estimate describes how weeks to arrest unfolds across the sample: at each observed moment, the survival curve shows the share of people who have not yet experienced the event. It handles right-censoring correctly (people whose event hasn't happened by the end of follow-up are treated as still at risk, not as non-events). Read the median and per-group event counts in the table as the anchor points; the curve itself reveals the **shape** of the timing (early drops vs. late attrition) that a single number cannot convey.

How age and prio predict weeks to arrest. The Cox proportional-hazards model finds that age and prio carry information about when the event happens: the overall likelihood-ratio test is significant, and the

model's discrimination is modest (the model discriminates better than chance, but not by a wide margin). Each row of the hazard-ratio table shows one predictor's effect — $HR > 1$ means the event happens sooner as that predictor increases (or is present, for a categorical level); $HR < 1$ means it happens later. Read the 95% CI alongside the point estimate: a CI that crosses 1 means that predictor's effect is uncertain in this sample. The proportional-hazards assumption was not fully met, so the hazard ratios should be read as average-across-time effects rather than constant multipliers; a stratified or time-varying model would refine that.

The findings present a nuanced picture of the factors influencing the length of time individuals remain free before a potential arrest following release. The primary focus of this evaluation centered on understanding how age, prior criminal history (prior offenses), and the provision of financial assistance relate to the timeline of re-arrest. When examining personal history, the data clearly show that an individual's age and their number of prior offenses are strongly connected to how quickly re-arrest occurs. Statistically significant findings mean that these patterns are highly unlikely to be random coincidence, indicating that age and prior criminal records provide meaningful predictive insight into the recidivism timeline. On the other hand, the investigation into whether financial aid extends the number of weeks an individual remains arrest-free yielded borderline or inconclusive results. While survival analysis and hazard modeling were used to trace the path over time between those who received aid and those who did not, the difference between the two groups sat precisely on the traditional threshold of statistical significance. Consequently, we cannot firmly conclude that financial aid directly alters the timeline to arrest in this sample. Rather than proving that financial aid has no impact whatsoever, this pattern indicates an absence of definitive evidence within this specific group of participants.

Conclusions

In summary, the strongest and most dependable conclusion supported by the evidence is that demographic and historical factors—specifically an individual's age and their count of prior offenses—are meaningful indicators of the timing of re-arrest. Programs or policy risk-assessment tools can reasonably rely on age and prior record as key components

when evaluating re-arrest timelines. Conversely, the evidence regarding the protective effect of financial assistance remains uncertain. Because the comparisons between financial aid recipients and non-recipients hovering right at the margin of statistical reliability, it would be premature to claim that financial aid is either entirely ineffective or definitively beneficial based solely on this dataset. Maintaining this careful distinction prevents over-interpreting weak trends while ensuring that genuine predictive relationships guide future decision-making.

Recommendations

Based on these observations, several actionable steps are recommended. First, because age and prior offenses demonstrated a clear, reliable link to the re-arrest timeline, these factors should continue to be monitored and validated in independent, prospective datasets to ensure the pattern holds across different locations and time periods. To better understand if interventions can directly alter these outcomes, future research should consider longitudinal tracking or experimental designs. Second, given that financial aid showed a borderline outcome, decision-makers should not abandon the concept of financial support entirely; instead, this relationship should be re-evaluated using a larger sample size to provide greater statistical power to detect smaller or subtle benefits. Finally, operational teams should ensure strict data collection protocols—eliminating duplicate entries and standardizing record-keeping—and consistently evaluate the practical magnitude of each finding alongside statistical tests to ensure policies are driven by meaningful real-world impacts.

Limitations

Several important constraints must be kept in mind when interpreting these findings. First, the underlying dataset contained duplicate records that required addressing, highlighting potential inconsistencies in data recording. Second, because this was an observational study rather than a controlled experiment, the reliable links observed between age, prior offenses, and time to arrest show association rather than direct cause-and-effect; unmeasured background factors, such as employment opportunities or local law enforcement practices, could influence these

relationships. Lastly, these results reflect a specific group of individuals over a observed period, meaning that extending these conclusions to broader populations or different geographic regions requires caution until further representative studies are conducted.

References

1. Cohen, J. (1988). Statistical Power Analysis for the Behavioral Sciences (2nd ed.). Lawrence Erlbaum Associates.
2. Seabold, S., & Perktold, J. (2010). statsmodels: Econometric and statistical modeling with Python. Proceedings of the 9th Python in Science Conference, 92–96.
3. Virtanen, P., Gommers, R., Oliphant, T. E., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods, 17, 261–272.

Python Code Appendix

Descriptive Statistics

```
df[[]].describe()
```

Frequency & Percentage Distribution

```
df['arrested'].value_counts(normalize=True)
```

Numeric-claim validation: PASSED (27 numbers checked). Engine v0.1.0 · Registry v0.2.0. Prose source: llm:gemin.