

Knife Crime Prevention Orders (KCPO): Impact Evaluation

A report by the MPS Strategic Insight Unit, April 2023

Executive Summary

The MPS's Strategic Insight Unit (SIU) has conducted an analysis to assess the impact of Knife Crime Prevention Orders on re-offending rates.

- This project was implemented as a randomised controlled trial, whereby knife crime offenders were randomly assigned either to receive a KCPO (the treatment group) or not (the control group).
- For a range of reasons (see accompanying qualitative report for more details), fewer people than expected were assigned to receive a KCPO. Additionally, very few of the individuals who were assigned to receive a KCPO actually received one. As such, we have very few observations in our analysis, which has limited our capability to distinguish impacts across treatment and control.
- Furthermore, we found that younger people (minors, and people up to 25 years old) were more likely to receive a KCPO. As such, the group of individuals who did receive a KCPO is not a random or representative sample of treatment group offenders.
- Descriptively, both people receiving KCPOs and people in the 'control' group have very high re-offending rates. They also re-offend with similar delays. We did not find any statistically significant differences in their rates of re-offending. Further analyses could consider the specific nature of each reoffending incident, and isolate whether knives were involved; however, the interpretation of findings from these analyses will likely be limited by the same challenges identified above.
- Due to possible selection into treatment, and to having a low number of observations the results need to be interpreted with care and may not reveal real causal effects. The lack of significance could indicate differences in the treatment and control, a lack of statistical capability to detect results, or a combination of the above.

1. Introduction

Knife crimes have been on the rise across England and Wales since 2014 (Home Office, 2023). The problem is particularly stark in London, where knife crime rates are three times higher than in the rest of England and Wales.

In response to the increasing problem of rises in knife-carrying behaviour and violent offences involving knives and sharp instruments, the Offensive Weapons Act of 2019 introduced several preventative measures designed to decrease re-offending in known knife crime suspects. Among these measures, the Act introduced Knife Crime Prevention Orders (KCPOs), a type of court order designed to prevent individuals from carrying knives or bladed articles in public places.

KCPOs would be assigned to individuals believed to be at risk of committing knife crimes, and could be tailored to each person. The KCPOs impose enforceable conditions on individuals, which could affect their incentives to commit future knife crime offences. The final decision on whether to pursue a KCPO is taken by the Crown Prosecution Service (CPS), and the order is ultimately awarded by the judge assigned to each case, who may or may not take into account the recommendations provided by the police officers involved in the case.

In this report, we evaluate re-offending rates for individuals who received a KCPO when compared to a control group of individuals who did not. The research design was originally planned to be a Randomised Controlled Trial (RCT). However, there were issues in the implementation which mean the treatment and control group are not perfectly comparable, and that we have a very small number of observations that are usable for this analysis. As such, the relationships described in the report cannot be interpreted as causal.

The first issue in the randomisation was a very low take up rate of KCPOs in the treatment group, meaning that those who were placed on a path which should have led to a KCPO often failed to reach the end of this journey. Many people who were assigned to receive a KCPO did not receive one, due to a long list of reasons which are described in the following section. From an initial 1,449 offenders who would be either in the 'treatment' or the 'control' groups of the experiment, only 220 ended up being analysed: 100 who were 'treated' (meaning they received a KCPO from CPS), and 120 who were in the 'control group'. We describe the implementation and how this affected the sample size in Section 2. As a result, we have lower statistical power (fewer observations limits our ability to distinguish effects in the treatment and control groups). This also calls for caution in interpreting the analysis, for the lack of significance in our effects may be due to not having sufficient observations.

Secondly, the assignment of an individual to the treatment group required multi-agency collaboration, and relied on high engagement from local police assigned to each case. However, the ultimate decision regarding whether to grant a KCPO or not fell on the judge assigned to the case, as KCPOs are made at the discretion of the court. People who were assigned to treatment might not receive a KCPO because a judge didn't deem it appropriate. In some cases, the individual circumstances of each case might have meant that a KCPO was not the best course of action. In yet another small number of cases, some people were not found guilty. However, there are other explanations for the lack of uptake besides some 'natural' attrition. The

report that accompanies this work highlights that there were differences in interpreting KCPOs from different judges, and some opposition to their implementation in some Courts.¹ In fact, the majority of individuals who were assigned to receive a KCPO didn't receive one. Our analysis reveals that age predicts the probability of offenders receiving a KCPO, with younger people being more likely to reach this milestone. This suggests that the treatment and control groups are somewhat different, and as such that there is non-random selection into treatment. In any RCT evaluation, the underlying assumption is that decision to provide treatment is uncorrelated with people's outcomes. In our (imperfect) setting, however, we cannot distinguish whether the differences observed in treatment and control group are due to receiving treatment or due to these systematic differences in the composition of the two groups. We also discuss these issues in the coming section.

There were no significant differences across treatment and control groups in terms of probabilities of re-offending, probabilities of being charged with an offence, and the number of re-offences committed. People in the treatment group and control group also appear to re-offend with similar delays. As such, any differences in average outcomes can only be read as illustrative and shall be interpreted with caution.

The rest of the report is organised as follows. Issues in the implementation of the trial and how that affects our sample size is explained in Section 2 and may prove useful for the design and implementation of future trials in similar contexts. Section 3 describes our selected outcome variables and crime definitions as well as the research methodology. Section 4 presents our results and a discussion, and Section 5 concludes.

2. Research Design and Issues in the Implementation

The evaluation of KCPOs was designed as a Randomised Controlled Trial (RCT) so as to provide comparability across treatment and control groups. The objective was that any differential effects could be interpretable as causal. However, the implementation was characterised by a very low take-up rate in KCPOs, which poses very important challenges for causal inference due to low statistical power, and selection. These issues are discussed in greater depth in the accompanying qualitative report.

There was an initial sample of 1,449 individuals who would be allocated to either treatment or control group. Initially the randomisation was designed to be 50/50. This meant that 50% of all individuals identified as eligible to partake in the experiment (i.e., suspects of knife crime) would be in the treatment group and 50% would be in the control group.

As the experiment progressed, the team observed that the randomly-assigned intention to deliver KCPOs to the treatment group was seldom realised in practice. Very few of the individuals initially assigned to receive a KCPO would ultimately

¹ For instance, some Courts felt that there were already other means in place to achieve falls in knife-carrying behaviour. See accompanying process evaluation.

receive one in court. In an attempt to compensate for this unexpected situation, the randomisation for future individuals entering the experiment was adjusted, with 80% of all knife crime suspects being assigned to treatment and the rest to the comparison group.

The KCPO implementation required a multi-agency effort from MPS and the judiciary. Much of the low take up in the treatment group was due to low engagement from police officers, in spite of efforts from MPS to drive up officer engagement.² After an individual had been assigned to the KCPO treatment, the police officer in charge of the case would have to write a statement to CPS after consulting with the Youth Justice Service. As such, there was an additional bureaucratic burden for officers dealing with cases of people initially allocated to receive a KCPO, which sometimes couldn't be met if there were other operational demands. The report that accompanies this work includes various interviews with MPS staff involved in the implementation that reveal some of these challenges and the perceptions from officers.

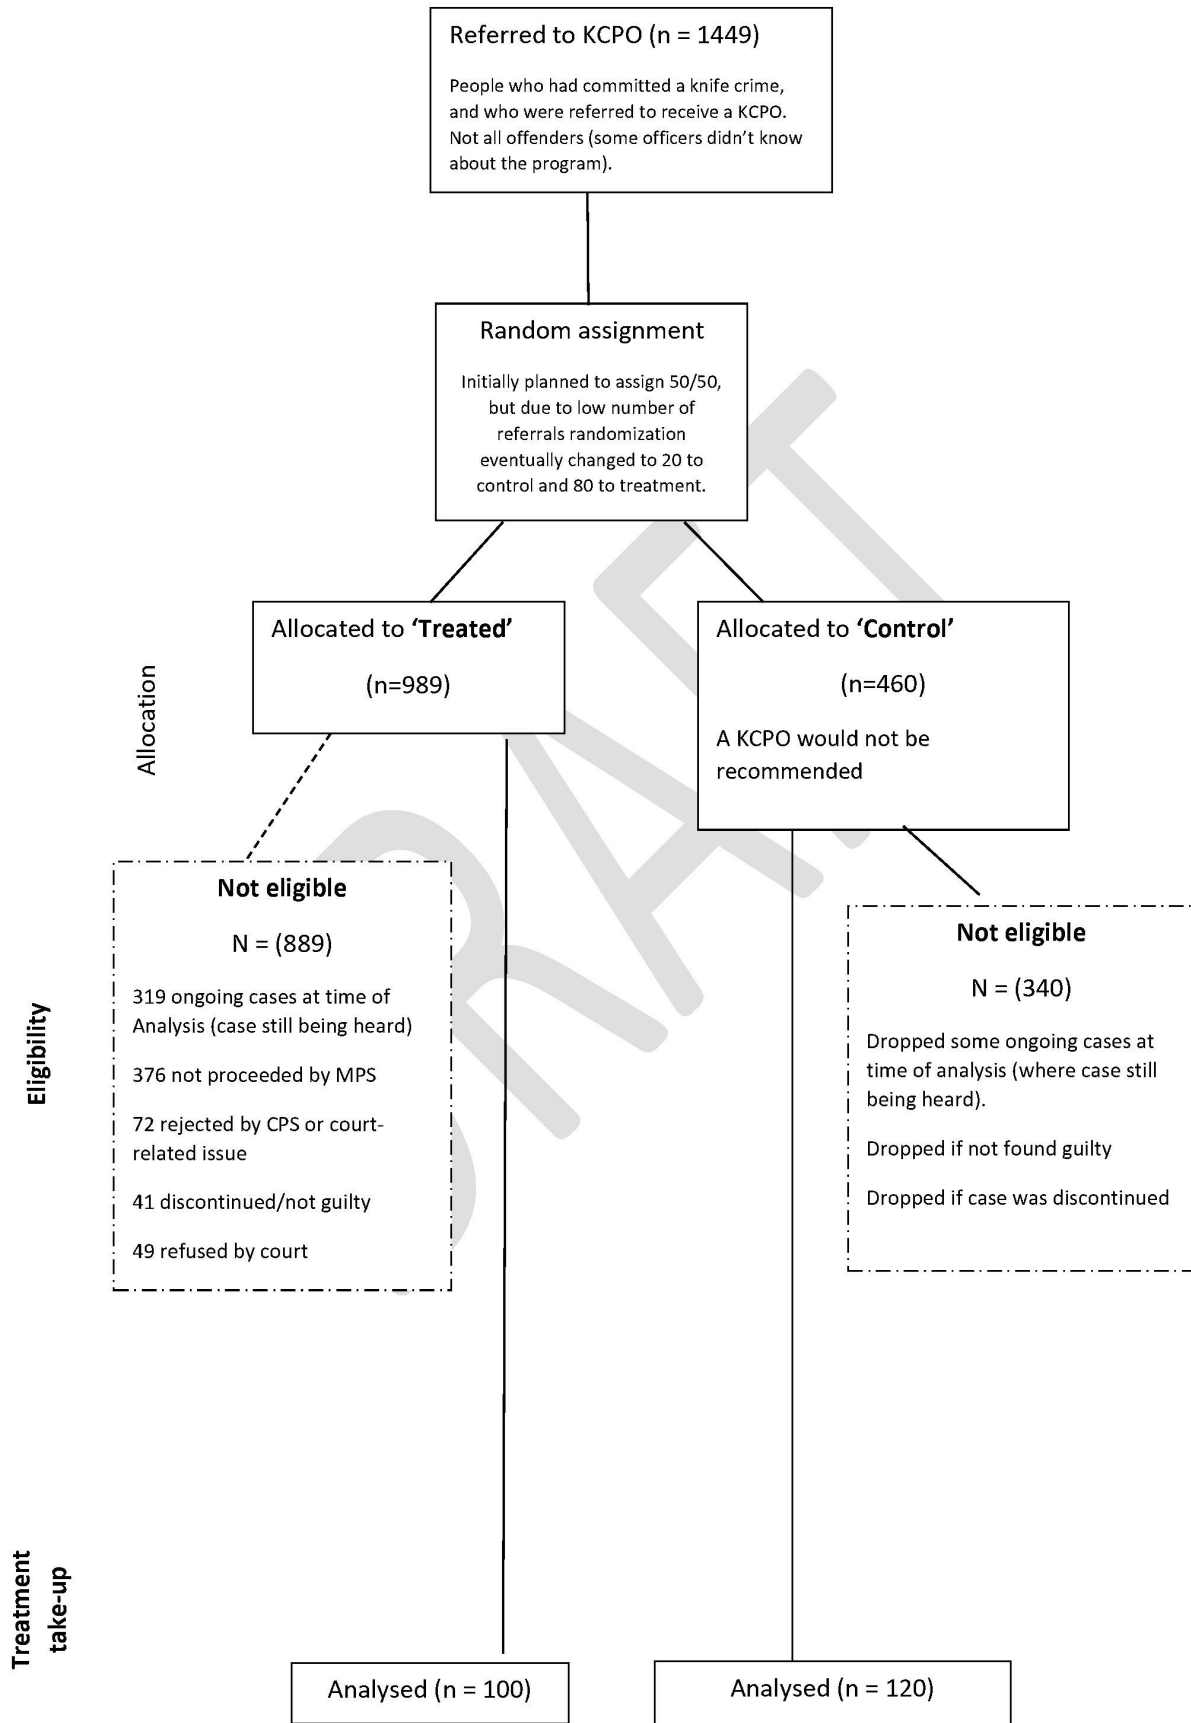
On the side of the judiciary, if a person pled guilty early-on in the trial process, the Court may not have read the recommendation from the police, and hence a KCPO would not be granted. The accompanying report also highlights misinterpretation from the judiciary as a main reason for low take up rates.³ Hence, some judges decided not to grant the KCPO even after being made aware of the recommendation, some KCPOs were refused by the courts, and yet another subset of cases were considered 'not guilty'. This limited the amount of people actually assigned to treatment to only 100.

On the control group side, many cases were dropped from this analysis if the investigation was still ongoing. In other cases, the observations were not deemed comparable; for instance, if an individual was not found 'guilty' or if the case was discontinued. This decision to exclude cases that were found not guilty at court was due to fact that KCPOs were only applied at the conclusion of the court process. That left our research team with 120 individuals in the control group.

² For instance, individual BCUs were contacted by senior police if they had low take-up rates.

³ Some interviewees highlighted a lack of familiarity with the legislation around KCPOS. See accompanying process evaluation.

Figure 1 – CONSORT diagram



3. Data and Methodology

Our analysis included cases for whom a KCPO was issued up to, and including, 3rd October 2022. We compiled re-offending rates from the CRIS (Crime recording information System) and COPA (Case Overview and Preparation Application used by MPS and CPS) records, which include details on offence type and offence date, together with some demographic information for offences taking place within the London Metropolitan Police force. We included data on reoffending recorded up to, and including, 21st February 2023.

There are 220 individuals in our study, 100 in the treatment group, and 120 in the control group. The majority of our sample are men, with only 5 women being included in the study. Across the two groups, they appear to be reasonably comparable in terms of gender, and ethnicity, but we observe some differences across age groups. Most people in the treatment group (82%) are under 25 compared to only 67% in the comparison group. This difference is not statistically significant, although it would be significant at a less-stringent 90% confidence level ($t=1.695$, $df=218$, $p=0.091$).

This may reveal some selection by various decision-makers, or the fact that youth cases (involving individuals below 25 years old) progressed through the courts far quicker. As such, who would obtain a KCPO depends on additional factors beyond the randomisation which may correlate with recidivism patterns. This implies comparisons might be biased, and calls for caution interpreting the results.

Table 1: Summary statistics, treatment vs control groups

Variable	Treatment (N=100)	Control (N=120)
Gender		
Male	98%	97.5%
Female	2%	3%
Ethnicity		
White	30%	31%
Black	60%	59%
Other	10%	10%
Age		
Mean age	22.4	25.0
Under 18	44%	42%
18 to 25	38%	25%
Over 25	18%	33%

We evaluated the differences in the treatment and control groups between receiving the Court order and the time of the analysis. The average number of months we observe a person since they are assessed in Court and either receive a KCPO (or not) is 134 days minimum and 572 days maximum, and the time differences between the two group are not statistically significant.

We assess four re-offending outcomes: the probability of being arrested, the probability of being charged, the number of offences being charged with, and the mean harm of offences as measured through the Cambridge Crime Harm Index (CCHI), which provides scores for different offences based on sentencing guidelines. We also evaluated the timing of reoffending in a survival analysis.

The most common re-offences were possession of weapons and violence. Possession of weapons was identified in 45% of all re-charges committed by people in the KCPO group, and 54% of all re-charges committed by people in the control group. A Table with the breakdown by offence type is included in the Appendix.

In the present report, we don't focus on knife crimes in isolation due to data limitations: it is often unclear whether an offence type involved a knife or not in the COPA and CRIS records.⁴ We also exclude domestic violence related offences, for they are outside of the scope of the KCPO objectives, and we exclude breaches to KCPOs as they can only affect the treatment group.⁵ Later analyses, however, should be able to look at the specific nature of each reoffending incident, and isolate whether knives were involved.

4. Results

We describe differences in means across the treatment and control groups in our four outcomes of interest, for which summary statistics are available in Table 1. As above-mentioned, any differences in outcomes are not easily attributable to having received a KCPO due to the lack of randomisation. The differences we measure could be driven, or biased by the non-random selection of treatment group cases which ultimately received a KCPO. Furthermore, the low number of observations limits our statistical power, making it difficult to detect statistical significance.

Each outcome is measured from the moment of the Court decision, which would have given a KCPO to people in the treatment group, and another decision to people found guilty of the crime in the control group. Since this is the moment in which the consequences of their crimes are revealed to offenders, we understand this would be the moment since the KCPO or another method of punishment would affect people's decisions to re-offend.

If this analysis was instead focused on the full sample of the RCT, the outcome measures would be based on a different approach. We would instead begin measuring re-offending from the moment of random assignment, and not from the moment of the court decision. In this case, however, the moment of random assignment occurred many months prior to court disposition. During this period, it seems likely that this randomly-assigned decision had very little impact on the offenders who were participating in the research. There were few (if any) indications

⁴ Offences recorded as "Possession of an offensive weapon" don't clarify whether the offence involved a knife or other types of weapons such as guns, hammers, baseball bats and knuckledusters. Offences recorded as "robbery" or "attempted murder" don't specify if a weapon was used during the commission of the offence.

⁵ There were 16 breaches of KCPOs within the treatment group.

to these offenders that their cases were proceeding in different directions. The first real impact upon them was largely limited to time after sentence was imposed.

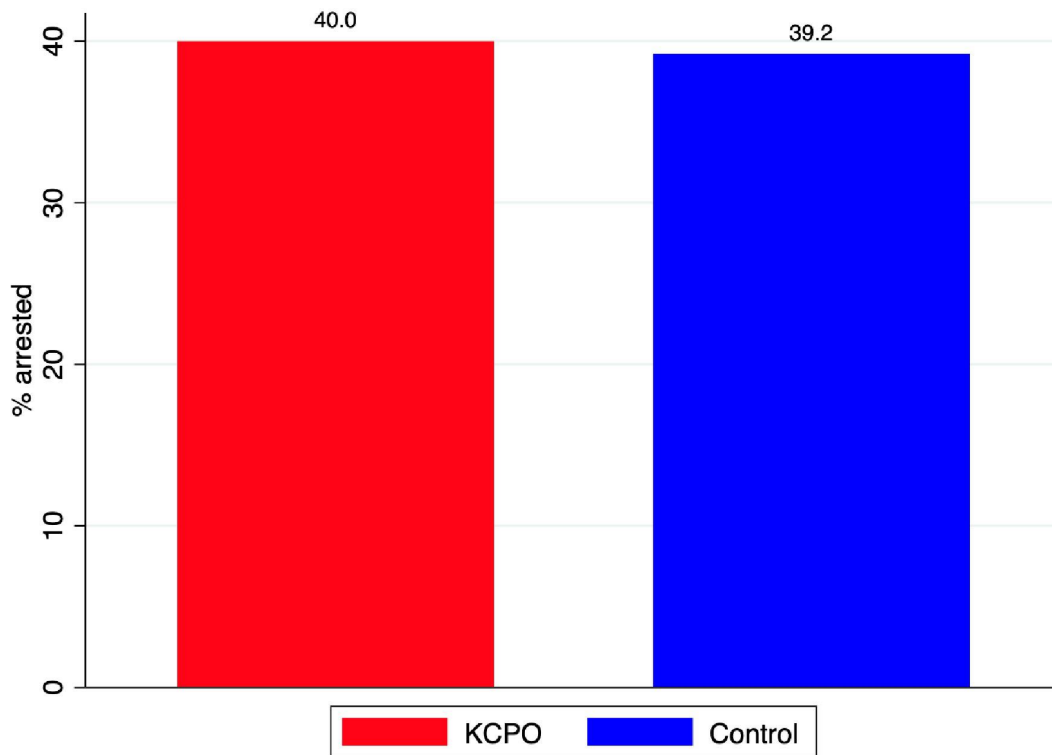
Even more importantly, our sub-sample for this report is limited to only the small minority which reached a final resolution in court. This sub-sample is uniquely tuned to this specific moment along the path of the criminal cases against them. In our judgement, this moment therefore presents the best available starting point for tracking further offending, at least for the purposes of this specific analysis.

The probability of being arrested after conviction is 40.0% for individuals in the KCPO group, and 39.2% in the comparison group; this difference is not statistically significant ($t = 0.00$, $df = 210.84$, $p = 1.000$).⁶

DRAFT

⁶ We also assessed the differences non-parametrically, finding similarly that there are no significant differences across treatment and control groups.


Figure 1: Probability of being arrested, KCPO vs comparison



Notes: the graph shows the probability of being arrested in the treatment and control groups from CRIS and COPA records. Offences related to domestic abuse and breaches to KCPOs are excluded from the analysis.

In terms of the probability of being charged, the average probability after receiving a KCPO is 29.0% for the treatment group, noticeably higher than the 24.2% for individuals in the control group. The difference is also not statistically significant ($t = -0.660$, $df = 206.99$, $p = 0.509$).


Figure 2: Probability of being charged



Notes: the graph shows the probability of being charged in the treatment and control groups from CRIS and COPA records. Offences related to domestic abuse and breaches to KCPOs are excluded from the analysis.

Our third outcome of interest is the average number of re-offences per offender. In the KCPO group, the mean number of charges is 0.97, whereas this is 0.86 for the comparison group. These modest differences are driven by just 3 offenders in the treatment group, each of whom present with more than 11 offences (i.e., the distribution of re-offences has a longer right tail in the treatment group). In the comparison group, the highest number of charges for a single offender was nine offences. The difference is not statistically significant ($t = -0.352$, $df = 176.55$, $p = 0.724$).


Figure 3: Mean number of charges per offender



Notes: the graph shows the average number of charges per offender in the treatment and control groups from CRIS and COPA records. Offences related to domestic abuse and breaches to KCPOs are excluded from the analysis.

We also compare the harm caused by individuals in the treatment and comparison groups. We use the Cambridge Crime Index, a widely used measure to assess harm of crime within the UK which uses the starting point from sentencing guidelines to rank harm from crimes based on the number of days in prison for that crime (Sherman, Neyroud and Neyroud, 2016). In the KCPO group, the average sum of crime harm is 340.5, whereas in the control this was 353.1. The difference in means suggests the comparison group were slightly more harmful than the KCPO group, however this difference is not statistically significant ($t = 0.074$, $df = 211.1$, $p = 0.941$).

Figure 4: Mean harm per offender

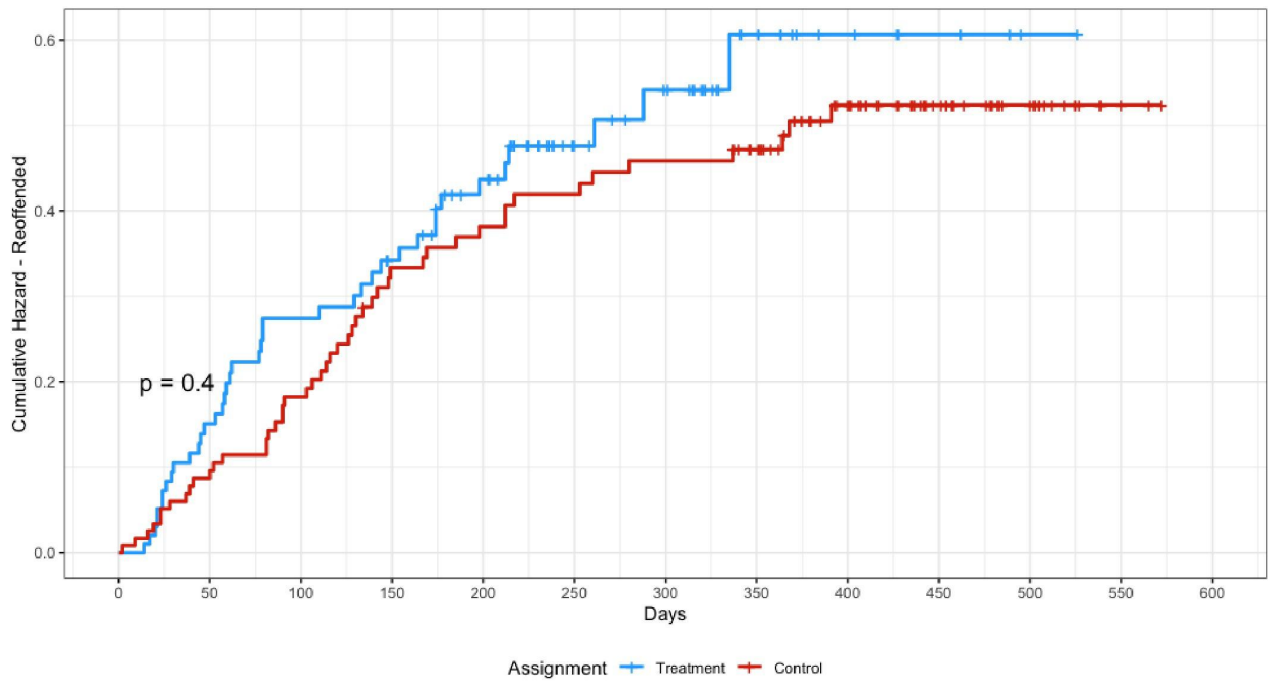


Notes: the graph shows the average crime harm per offender in the treatment and control groups from CRIS and COPA records. Offences related to domestic abuse and breaches to KCPOs are excluded from the analysis. We use the Cambridge Crime Harm Index, which is a continuous measure ranging between 1 and 5475, with scores based on number of days in prison determined from the starting point in sentencing guidelines.

Finally, we assess the impact of KCPOs on the timing of re-offending leveraging a survival analysis. Our study evaluates whether KCPOs delay the time to first re-arrest compared to no KCPO being issued. To do so we assess how long it took people to re-offend after the Court decision, which is the time in which punishment is observed by individuals.

The plot in Figure 5 shows the proportion of people in the sample who had re-offended by a given date after the Court order was granted. The interpretation is as follows: 100 days after the court decision, 25% of people in the KCPO group had been re-arrested, compared to about 20% in the control group. By 150 days after the court decision, the two groups show similar probabilities of re-arrest. The differences are nonetheless not statistically significant.

Figure 5: Cumulative hazard of reoffending for treatment and control



Notes: the graph shows the probability of re-offending by date, for people in the treatment and control groups. Offences related to domestic abuse and breaches to KCPOs are excluded from the analysis.

We cannot conclude from this analysis whether the KCPOs have an impact on recidivism. The observed similar rates of reoffending could be due to a) people in the KCPO group being on average younger, which may affect their engagement in crime; b) lack of statistical power to assert differences between observations; c) a combination of the above.

5. Conclusions

Our analysis highlights first, the difficulties in RCT design in the context of multi-agency interventions. In this case, receiving a KCPO required the long-term engagement of local police officers, and agreement by the CPS and from the judiciary. This full alignment of conditions was often not achieved. Overall there was extremely low treatment integrity with currently only 10.1% assigned to receive a KCPO given one. This percentage may increase slightly as cases progress. However, even if every progressing case receives a KCPO this will achieve a maximum of 42.3%. This is still well below the standard of treatment integrity necessary for high internal validity, which would provide increased confidence in the ability to draw causal conclusions from the results.

There also was very high attrition in both randomly-assigned groups, and only 15% of all the individuals who had been assigned to either treatment or control were analysed. In the treatment group, people who had been assigned to receive a KCPO would only receive one if the police officer assigned to the case went through the bureaucratic burden of completing the necessary forms, if the judge read the recommendation from the police, and ultimately the judge determined a KCPO was the optimal solution for each individual. Several individuals were not analysed if the cases were dropped, discontinued, not found guilty of a knife crime offence, or if the case was ongoing at the time of the analysis.

In the control group, only people deemed reasonably 'comparable' to the remaining treatment participants were analysed. People whose case was dropped, discontinued, who had an ongoing case at the time of the analysis, or who were not found guilty were dropped. The process of defining members of the control group as "comparable" unavoidably includes some rather subjective decision rules, and alternate rules could either increase or decrease the size of the comparison group. Nevertheless, the stark absence of difference between the two groups suggests that any changes to the comparison group would produce only marginal changes to these findings.

It is also likely that the characteristics of those who did (and did not) reach these thresholds will correlate with their probabilities of re-offending. Our summary statistics show that younger people (below 25 years old) were slightly more likely to receive KCPOs than those who were older. There are quite likely to be other variables, including ones that we cannot examine here, which also correlate with both the delivery of KCPOs and later re-offending. Therefore, the realisation of random assignment was hampered in this experiment. Even if randomisation had been robust, the low statistical power of this analysis makes it very difficult to measure causal differences between the treatment and control groups.

Future RCTs should approach interventions like KCPOs differently. Firstly, the random assignment should take place much later in the lifespan of a criminal investigation, as close to the tangible delivery of the intervention as possible. Secondly, eligibility for random assignment should be more tightly constrained, and eliminate those whose cases that are unlikely to reach the court resolution needed to deliver the treatment to those assigned to receive it.

Finally, such research should carefully assess how many parties are involved in treatment assignment, understand whether different institutions involved in treatment assignment have different incentives to participate or not, and study whether all the organisations involved know the specifics of the intervention. By leaving less aspects open to interpretation, and minimising the number of people involved one is likely to reduce attrition significantly.

We also strongly recommend running a short pilot trial ahead of the full RCT, which may help organisations understand barriers to implementing the intervention they wish to test, and allow randomisation to be set at the most appropriate point for ensuring high treatment integrity.

Despite these limitations, we compared outcomes for people who received a KCPO to people who did not, and have presented summary statistics describing their re-offending rates.

Both the KCPO and the control group are at a fairly high risk of reoffending with average re-offending rates of 40% between 4 and 19 months later. In comparison, in the whole population of offenders the re-offending rate is 24.3% after 12 months post-arrest (Ministry of Justice, 2023). There are small differences in terms of re-offending frequency, with the KCPO group having a small number of more highly frequent offenders. However, the control group contains some higher-harm offences (slightly fewer offences, but more harmful). There were no statistically significant differences in terms of the delays in re-offending, but the KCPO group appears to re-offend slightly (but not significantly) faster.

From these results, we cannot conclude whether KCPOs have a significant impact on people's reoffending rates or not.

References

Home Office (2023) Knife crime statistics, available via <https://commonslibrary.parliament.uk/research-briefings/sn04304/>, last visit 25th April 2023.

Ministry of Justice (2023) Proven reoffending statistics: January to March 2021, available via <https://www.gov.uk/government/statistics/proven-reoffending-statistics-january-to-march-2021/proven-reoffending-statistics-january-to-march-2021>, last visit 26th April 2023.

Sherman, L., Neyroud, P. W., & Neyroud, E. (2016). The Cambridge crime harm index: Measuring total harm from crime based on sentencing guidelines. *Policing: a journal of policy and practice*, 10(3), 171-183.

Appendix

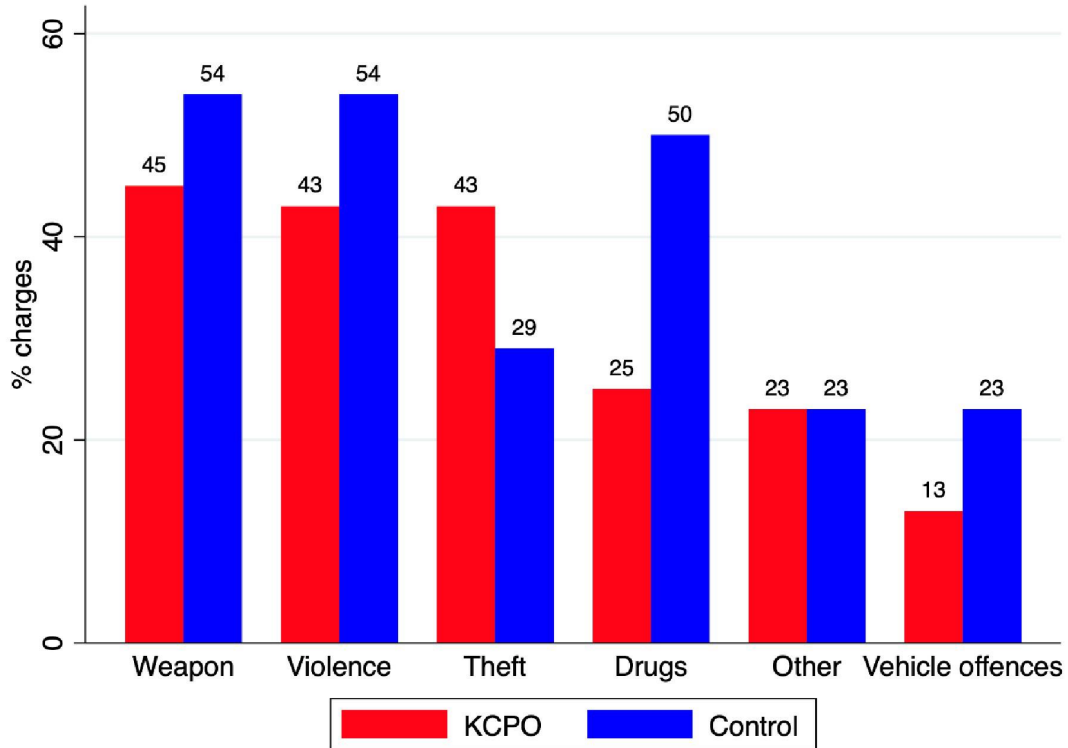
Table A1: Age distribution in treatment vs control group

Age	Treatment (N=100)	Control (N=120)
-----	----------------------	--------------------

Range	14 – 63	14 – 56
Mean	22.4	25
Median	18	18

Notes: the graph shows the age distribution at the time of treatment assignment for people in the treatment and control groups.

Figure A1: Re-Offending crime charges by crime type



Notes: Types of offences charged that KCPO and control group offenders were arrested for after random assignment. Offenders may be charged for multiple offences hence why the total percentages add up to greater than 100%.