



This document provides answers to frequently asked questions about Justice Data Lab (JDL) analyses and reports, to help customers and other interested parties to understand the background to the analyses and how to interpret the results.

Contents

What is the JDL and how does it work?	2
What is the purpose of a JDL analysis?	2
How is reoffending measured?.....	2
How is the impact of an intervention analysed?	3
What caveats should be considered when reading a JDL report?	3
Who is excluded from JDL analyses?	5
What measures of reoffending are used in the analysis?	5
How should the numbers in the report be interpreted?	6
Guide to JDL report symbols.....	7
What do the confidence intervals on the graphs mean?	7
What is the relationship between the measurements and the estimates?	8
Can the true impact of the intervention ever be outside the estimated range?	8
What are percentage points?	9
What is the meaning of statistical significance?	9
If the results of an analysis are non-significant, does it mean that the intervention has no impact?	9
How can a measure show a significant result when another measure shows a non-significant result for the same intervention?	10
Why do some JDL reports contain more than one analysis?	10
Where can I find previous JDL analyses and methodology documentation?	11
Contact points.....	12

What is the JDL and how does it work?

The Justice Data Lab is a small team from Analytical Services within the Ministry of Justice that supports organisations that provide offender services, by allowing them easy access to aggregate reoffending data specific to the group of people they have worked with. This service is intended to support organisations in understanding their effectiveness at reducing reoffending.

Participating organisations supply the Justice Data Lab with details of the offenders who they have worked with, and information about the services they have provided. The JDL matches these individuals to the reoffending datasets held within the Ministry of Justice and uses statistical modelling techniques to generate a matched comparison group of individuals with very similar characteristics. As a standard output, the Justice Data Lab supplies aggregate one year proven reoffending rates, frequency of reoffending and number of days to the first reoffence for the group of offenders the organisation has worked with, and those of the matched comparison group of similar offenders.

The reoffending outcomes for the organisation's group and the matched comparison group are also compared using statistical testing to assess the impact of the organisation's work on reducing reoffending. The results are then returned to the organisation with explanations of the key metrics, and any caveats and limitations necessary for interpretation of the results.

Finally, the tailored reports produced for each organisation are published on gov.uk to promote transparency and ensure that findings produced through this service can be used by others to improve the rehabilitation of offenders.

What is the purpose of a JDL analysis?

A JDL analysis adds to the evidence about the way in which an intervention with offenders affects their reoffending behaviour. This helps to determine whether the intervention reduces reoffending among its participants, and by how much. The analysis also examines the impact of the intervention on reoffending at different levels of severity.

The results of a JDL analysis can be used to estimate the impact that an intervention would have on people who are similar to those who received it. This is designed to help providers decide on whether to offer the intervention to more people, or whether to change aspects of it.

How is reoffending measured?

In order to be included in a JDL analysis, a person must previously have committed a criminal offence that resulted in a court conviction, caution, reprimand or warning in England or Wales. This is called their 'index offence'. Reoffending is recorded from the 'index date', which is the date on which the person left custody or received a non-custodial sentence, caution, reprimand or

warning as a result of the index offence. If there is more than one offence relating to the same index date, the most serious one is used as the index offence.

JDL analyses look at 'one-year proven reoffending', which records offences that the person committed during a one-year period starting on the index date and that resulted in a court conviction, caution, reprimand or warning in England or Wales during the same period or a further six-month waiting period.

If a person commits an offence more than one year after their index date, or if they are convicted of an offence more than eighteen months after their index date (or never convicted), this will not be recorded as a re-offence in a JDL analysis.

For a full description of the proven reoffending statistics used by the JDL, see pages 8-11 of: https://assets.publishing.service.gov.uk/media/6793d1611e5e5352cb24bf62/Technical_guide_to_proven_reoffending.pdf.

How is the impact of an intervention analysed?

A JDL analysis provides *measurements* of the reoffending behaviour of a 'treatment group', which is a group of people who received an intervention around the time of their index date or during their index prison sentence, and whose identifying details are provided by the organisation running the intervention. It also provides measurements of the reoffending behaviour of a larger 'matched comparison group' (sometimes called a 'control group'), which contains people who are similar to those in the treatment group but who have not received the intervention. The measured differences between the groups are used to make *estimates* of the impact of the intervention on the reoffending behaviour of the treatment group. The estimates also provide evidence of the impact that the intervention would have on the reoffending behaviour of any people who are similar to those in the treatment and comparison groups. This is useful when deciding whether to provide the intervention to more people.

The comparison group is expected to exhibit the reoffending behaviour that the treatment group would have exhibited if they had not received the intervention. The members of the comparison group are therefore chosen for their similarity to the members of the treatment group on a wide range of characteristics that are known to be generally related to offending behaviour, including demographics, employment history, criminal history and individual risks and needs. They are selected from a pool of almost all the offender records in England and Wales that have index dates in the same years as those of the treatment group.

What caveats should be considered when reading a JDL report?

When interpreting the evidence presented in a JDL report, the following caveats should be considered:

- The differences between the reoffending behaviour of the treatment and comparison groups could be due to unobserved factors as well as to the impact of the intervention. This is because the comparison group must be selected to match the people in the treatment group, and the matching can only take account of characteristics that have been observed in the data sets used. The matching process uses individual information about offenders in England and Wales, which is drawn from reliable administrative data sets. This covers a wide range of characteristics that are related to offending behaviour, including demographics, employment history, criminal history and individual risks and needs. However, it is still possible that the reoffending behaviour of the treatment and comparison groups will differ due to unobserved characteristics such as the impact of other interventions, motivation to change offending behaviour or the complexity of personal problems.
- Reliable results depend on good matching between the treatment and comparison groups. 'Standardised differences' are provided to indicate the quality of matching in an analysis, and are individually rated as 'good' (-5% to 5%), 'reasonable' (-10% to -5% or 5% to 10%) or 'poor' (below -10% or above 10%). Each standardised difference is a measure of the difference between the treatment and comparison group averages for one characteristic. The more standardised differences are rated as good, the more reliable the results of the analysis are. JDL analyses aim to achieve the best possible matching with the data available.
- It may not be possible for the treatment group to include everyone who has received the intervention, and the impact of the intervention on those who are included may be different to the impact on those who are not. The impact of the intervention could also be quite different on an entirely different group of people. For this reason, the reoffending behaviour measured in an analysis should not be directly compared either to the reoffending behaviour measured in any other analysis or to figures such as national averages.
- The impact of the intervention may differ for each person who receives it. A JDL analysis can only estimate the impact of the intervention on the reoffending behaviour of the treatment group as a whole.
- Some re-offences committed during a one-year period are not recorded in JDL analyses. Only proven reoffending is recorded, which is generally an underestimate of actual reoffending. In addition, a re-offence is not included if it is proven more than six months after the end of a person's one-year reoffending period. This means that a re-offence is more likely to be included if it is less serious, because the time between offence and conviction is generally shorter for less severe offences. Unrecorded offences affect the measured reoffending rates of both the treatment and comparison groups.

Who is excluded from JDL analyses?

Some of the individuals submitted by a customer may be excluded from the final analysis. They fall into the following groups:

- People whose details cannot be linked to the Police National Computer – this may be due to their identifying information being incorrect, or it may be because they have not been convicted of any offences.
- People whose details cannot be linked to the JDL's reoffending data – this may be because the dates submitted for their intervention do not correspond to the date of any recorded community sentence or release from custody, or it may be because their conviction or release is too recent to provide one year reoffending information at the time of analysis.
- People who cannot be matched to anyone in the comparison group. A fair comparison cannot be made for these people, because they are not similar enough to anyone in the JDL's records who has not received the intervention.
- People who reoffended before their intervention began – this applies to those whose intervention began after they had received a community sentence or been released from prison. They are excluded because they committed a reoffence before the intervention was able to make an impact on their behaviour.
- People who had committed a proven sexual offence at any time before their intervention began. Sex offenders are presently excluded from JDL analyses, because their reoffending patterns are generally very different from those of non-sex offenders. We are currently investigating the possibility of analysing interventions designed for sex offenders.
- People aged under 18 at the time of their community sentence or release from custody – they are excluded unless the intervention is designed for young offenders. This is because under-18s may receive different sentence types to over-18s.

What measures of reoffending are used in the analysis?

In a JDL report, measurements and estimates are presented for up to fourteen measures of one-year proven reoffending. Three are headline measures of overall reoffending, nine are measures of reoffending seriousness and two are measures of custodial sentencing. They are:

- *Reoffending rate* – the number of people who committed a proven re-offence, expressed as a percentage of the group
- *Reoffending frequency* – the number of proven re-offences committed, expressed per person

- *Average time to first re-offence* – the average number of days between a person's index date and the date on which they committed their first proven re-offence, including only those who re-offend
- *Indictable-only reoffending rate* – the number of people whose first proven reoffence must be tried at a Crown Court, expressed as a percentage of the reoffenders
- *Triable-either-way reoffending rate* – the number of people whose first proven reoffence may be tried either at a Crown Court or a magistrates' court, expressed as a percentage of the reoffenders
- *Summary reoffending rate* – the number of people whose first proven reoffence is usually tried at a magistrates' court, expressed as a percentage of the reoffenders
- *Indictable-only reoffending frequency* – the number of proven reoffences that must be tried at a Crown Court, expressed per reoffender
- *Triable-either-way reoffending frequency* – the number of proven reoffences that may be tried either at a Crown Court or a magistrates' court, expressed per reoffender
- *Summary reoffending frequency* – the number of proven reoffences that are usually tried at a magistrates' court, expressed per reoffender
- *Custody rate* – the number of people who received a custodial sentence for their first proven re-offence, expressed as a percentage of the re-offenders
- *Custody frequency* – the number of custodial sentences received as a result of reoffending, expressed per re-offender

The three overall measures are included in every report, and the others are included if there are enough people in each category to allow reliable estimates to be made. Offences are classified into three categories of court outcome: *indictable-only* offences are the most serious and must be tried at a Crown Court, *triable-either-way* offences may be tried at a Crown Court or a magistrates' court, and *summary* offences are usually tried at a magistrates' court. For more information on offence types in these categories, see page 36 of:

https://assets.publishing.service.gov.uk/media/6793d1611e5e5352cb24bf62/Technical_guide_to_proven_reoffending.pdf

For the latest guidance to court statistics, please see:

<https://www.gov.uk/government/publications/a-guide-to-criminal-court-statistics>

How should the numbers in the report be interpreted?

Summary boxes and graphs in JDL reports express some measures per 100 people instead of per person or as a percentage, in order to make the numbers more meaningful to the reader. For example, a reoffending rate of 40% and a reoffending frequency of 1.2 offences per person

in the treatment group may be written as: "For 100 typical people in the treatment group, 40 people committed a proven re-offence within a one-year period. They committed 120 proven re-offences during the year."

Totals may not appear to equal the sum of the component parts because numbers have been rounded.

For a full description of JDL methodology, see:

www.justice.gov.uk/downloads/justice-data-lab/justice-data-lab-methodology.pdf

and:

www.gov.uk/government/uploads/system/uploads/attachment_data/file/506327/methodology-review-response.pdf

Guide to JDL report symbols

The following symbols are used throughout the JDL reports:

Headline outcome measures:



Number of re-offenders in 1 year



Frequency of re-offences in 1 year



Time taken to re-offend

Indication of results:



Significant decrease



Significant decrease (used for days to first re-offence outcome only)



Significant increase



Significant increase (used for days to first re-offence outcome only)



Non-significant decrease



Non-significant increase

What do the confidence intervals on the graphs mean?

A confidence interval shows the estimated range of a reoffending measure for one group. For example, the reoffending rate of a treatment group could be 40% and the confidence interval could cover the range from 35% to 45%. This would mean that the treatment group had a reoffending rate of 40%, and that similar people who received the intervention would be expected to have a reoffending rate that is between 35% and 45%.

What is the relationship between the measurements and the estimates?

The measurements give reoffending information for the treatment and comparison groups used in the analysis. The reoffending behaviour that is measured depends on the impact of the intervention, but it also depends on random factors affecting the particular choice of people for those groups.

The estimates quantify the impact of the intervention, allowing for the random factors. Each estimate is given as a range of numbers, and the size of the impact is expected to be somewhere within this range.

For example: the measurement of the difference in the one-year reoffending frequency could be -1.5 offences per person, with the impact estimated to be between -2.0 and -1.0 offences per person. Over a period of one year, this would mean that the treatment group committed an average of 1.5 fewer re-offences per person than the comparison group, and that the intervention would be expected to prevent between 1.0 and 2.0 re-offences per person if it were provided to similar people in future.

In a JDL analysis, an estimated range is always centred on the measurement that it is based upon. In general, the range becomes narrower as the treatment group becomes larger. This means that the impact of an intervention can be estimated more precisely when the treatment group is large.

Can the true impact of the intervention ever be outside the estimated range?

Yes. The estimates are calculated using a standard method called '95% confidence'. This means that the underlying difference between the reoffending behaviour of the treatment and comparison groups is expected to be within the estimated range for 95% of JDL results: for another 2.5% of results the underlying difference will be lower than estimated, and for the other 2.5% of results it will be higher. It does not mean that the estimated range should be regarded as "95% likely to be correct" in every case, because each result should be judged individually using all the available evidence.

It is also possible that the underlying difference between the reoffending behaviour of the treatment and comparison groups may be caused by unobserved differences between the groups as well as by the impact of the intervention. The estimated range does not take account of any unobserved differences.

What are percentage points?

Percentage points (pp) are used to express the difference between two percentages. For example, the difference between 5% and 10% is 5 percentage points, and the difference between 85% and 90% is also 5 percentage points.

What is the meaning of statistical significance?

A statistically significant result means that random factors appear to be an unlikely explanation for the measured difference between the treatment and comparison groups. It is robust evidence that there is an underlying difference between the reoffending behaviour of the groups. For example, the estimated range for the difference in the reoffending rate could be -5 percentage points to +3 percentage points. This is not statistically significant because it appears plausible that the difference is zero. But if the estimated range is -10 percentage points to -2 percentage points, this is a statistically significant result because it appears unlikely that the difference is zero.

Statistical significance is a guideline. It acts as a flag to highlight the most convincing pieces of evidence. For more guidance on understanding statistical significance, see: www.thinknpc.org/publications/understanding-statistical-significance/npc_understanding-statistical-significance_final

If the results of an analysis are non-significant, does it mean that the intervention has no impact?

No. A non-significant result means that it is plausible that the intervention has a positive impact, a negative impact or no impact on reoffending, based on the JDL analysis alone. For example, the estimated range for the difference in the reoffending rate could be -5 percentage points to +3 percentage points. It would then be plausible that the impact was -5 percentage points, or +3 percentage points, or anything in between. The only way to get robust evidence of the direction of the impact would be to include more people in the analysis.

A non-significant result provides useful information. In the example above, it appears unlikely that the impact is greater than -5 percentage points in the negative direction or greater than +3 percentage points in the positive direction.

Every measurement provides evidence about an intervention, even if the result is non-significant. The measurement is at the centre of the estimated range – so, in the example above, the measured difference in the reoffending rate would be -1 percentage point. This means that the most plausible conclusion is that the intervention reduces the reoffending rate by 1 percentage point. It is a promising result, but the lack of statistical significance means that the conclusion is not very robust.

Results from a JDL analysis should always be placed into the wider context of offender rehabilitation. Each result provides a piece of evidence about the effectiveness of an intervention. Other important pieces of evidence include the methodology of the intervention and reliable accounts of its impact on specific individuals. Effectiveness can also be demonstrated through outcomes other than reoffending behaviour, such as health or quality of relationships, which may lead to a reduction in reoffending after a period of more than one year. The overall picture of an intervention is made up of all these pieces. A JDL report is designed to provide reliable, quantitative evidence with which to judge the impact of the intervention in the specific area of reoffending behaviour over a one-year period.

How can a measure show a significant result when another measure shows a non-significant result for the same intervention?

A statistically significant result occurs when a difference in reoffending behaviour between the treatment and comparison groups becomes apparent from its estimated range. Each measure records a different aspect of reoffending behaviour, and the impact of the intervention on each aspect can be different. This means that the impact on one measure may be large enough to become apparent, while the impact on another measure in the same analysis may not be large enough. It does not necessarily mean that the impacts act in opposite directions.

For example, the estimated impact of an intervention on the reoffending rate could be between -10 percentage points and +1 percentage point (non-significant), while the estimated impact on the reoffending frequency could be between -1.5 and -0.5 offences per person (significant). It appears likely that the intervention reduces both the rate and frequency of reoffending, but there is a small chance that it reduces the frequency without reducing the rate.

Why do some JDL reports contain more than one analysis?

For interventions that are not gender-specific, each JDL report contains analyses for males and females separately. This is due to the known differences in reoffending behaviour.

Each JDL report contains 'national' and 'regional' analyses if the intervention takes place in a specific geographical area. All of these look at almost the same treatment group, but each has a different matched comparison group. Multiple analyses broaden the comparison between people who have received an intervention and people who have not, providing more detailed and robust evidence of the impact of the intervention.

All analyses use characteristics such as demographics, employment history and criminal history when matching the comparison group to the treatment group, as well as criminogenic risks and needs. The risks and needs information comes from the Offender Assessment System (OASys), which records data from individual offenders on factors such as accommodation, education and

work skills, substance misuse, mental health, relationships and attitudes towards offending. These factors can be important in matching a comparison group that is similar to the treatment group in terms of offending behaviour, but they may not be available for everyone in the treatment and comparison groups. For a full description of the integration of OASys data into JDL analyses, see: www.gov.uk/government/uploads/system/uploads/attachment_data/file/491688/oasys-methodology.pdf

National analyses use a comparison group of people who have been convicted of offences anywhere in England or Wales. If an intervention takes place in a specific geographical area, a regional analysis is also conducted using a comparison group of people who have been convicted in the same area. The national analysis selects the comparison group from the largest possible pool of people, while the regional analysis avoids any unobserved differences in offending behaviour between people in different areas.

Where can I find previous JDL analyses and methodology documentation?

An Excel workbook is published with every JDL release – this has a cover sheet with links to previous releases and individual reports. As well as JDL analyses, previous publications have also included review papers synthesising JDL results to date, feedback following the 2-year pilot phase of the JDL service and a number of methodology papers. These can all be found in the ‘Documentation’ tab of the same Excel workbook.

The JDL have produced a Power BI 'What Works Tool'. This is a visual presentation of the results of the analyses conducted by the JDL. Users can: observe the summary statistics of past analyses, group and filter past analyses by different categories, and view summary statistics in graph and table format.

Please note that developments to datasets used by the JDL may occur, and the JDL adopt any changes to variables to be consistent across statistics produced by the MoJ.

Contact points

Press enquiries should be directed to the Ministry of Justice press office:

<https://www.gov.uk/government/organisations/ministry-of-justice/about/media-enquiries>

Other enquiries about the analysis should be directed to:

Other enquiries about the analysis should be directed to:

Justice Data Lab team

Email: justice.datalab@justice.gov.uk

General enquiries about the statistical work of the Ministry of Justice can be e-mailed to:

statistics.enquiries@justice.gsi.gov.uk.

General information about the official statistics system of the United Kingdom is available from

<https://gss.civilservice.gov.uk/policy-store/code-of-practice-for-statistics/>

© Crown copyright

Produced by the Ministry of Justice

Alternative formats are available on request from justice.datalab@justice.gov.uk