

[Print Tutorial](#)

ACH Tutorial

Background

[What Is ACH?](#)

[When and Why to Use ACH?](#)

[How Is ACH Different?](#)

[Refuting vs. Confirming Hypotheses](#)

[ACH Is a 9-Step Process](#)

[ACH as a Collaborative Process](#)

[When ACH May Not Meet Your Needs](#)

Project Management

[Classification and Title](#)

[Adjust Column Width](#)

[Hide/Show Columns](#)

[Changing Colors & Font](#)

[Printing](#)

Building the Matrix

[Enter Hypotheses](#)

[Enter Evidence](#)

[Evidence Notes](#)

[Enter Type of Evidence](#)

[Enter Date or Time](#)

[Enter Code](#)

[Assess Consistency of Evidence](#)

[Assess Weight of Evidence](#)

[Credibility of Evidence](#)

[Relevance of Evidence](#)

[Flag Column](#)

Analytical Functions

[Sort Evidence](#)

[Create Duplicate Matrix](#)

[Partition the Evidence](#)

[Calculation of Scores](#)

[Show Graphics](#)

Drawing Conclusions

[Interpreting the Inconsistency Scores](#)

[Potential Pitfalls](#)

[Why No Probability Calculation?](#)

[Sensitivity Analysis](#)

[Identify Milestones](#)

[Reporting ACH Results](#)

Background

What Is ACH?

Analysis of Competing Hypotheses (ACH) is a simple model for how to think about a complex problem. It is an analytic process that identifies a complete set of alternative hypotheses, systematically evaluates data that is consistent and inconsistent with each hypothesis, and rejects hypotheses that contain too much inconsistent data.

ACH takes you through a process for making well-reasoned, analytical judgments. It is particularly useful for issues that require a careful weighing of alternative explanations of what has happened or is happening. ACH can also be used to provide early warning or help you evaluate alternative scenarios of what might happen in the future. ACH helps you overcome, or at least minimize, some of the cognitive limitations that make prescient intelligence analysis so difficult; it helps clarify why analysts are talking past one another and do not understand each other's interpretation of the data. ACH is grounded in basic insights from cognitive psychology, decision analysis, and the scientific method. It helps analysts protect themselves from avoidable error and improve their chances of making the right call.

This software provides a structured process for breaking a complex analytical problem down into its component parts - a full set of hypotheses (i.e., alternative explanations or outcomes), evidence and arguments that are useful in assessing these hypotheses, and judgments about the consistency or inconsistency of each item of evidence with each hypothesis. The software steps you through a process that helps you question your assumptions and gain a better understanding of the issue. The value of ACH is measured by the extent to which it helps you see an issue from alternative perspectives, prods you to look for additional evidence you had not realized was relevant, helps you question assumptions, identifies the most lucrative future areas of investigation, and generally stimulates systematic and creative thinking about the issue at hand.

The hypotheses, evidence, and analysis of the evidence are entered into a matrix that becomes a record of your thought process in analyzing a given topic or problem. This written record of your thought process is what helps you deal with the complexity inherent in most analytical problems. The software also allows you to sort and compare evidence in a variety of analytically-useful ways.

[Contents](#)

When and Why to Use ACH?

Use ACH when the judgment or decision is so important that you can't afford to be wrong. Use it to record and organize relevant evidence prior to making an analytical

judgment or decision. Use it to identify and then question assumptions that may be driving your thinking, perhaps without realizing it. Use it when the evidence you are working with may be influenced by denial and deception. Use it when gut feelings are not good enough, and you need a more systematic approach that raises questions you had not thought of. Use it to prevent being surprised by an unforeseen outcome. Use it when an issue is particularly controversial and you want to highlight the precise sources of disagreement. Use it to maintain a record of how and why you reached your conclusion.

[Contents](#)

How Is ACH Different?

The conventional intuitive approach to analysis is to focus on what we suspect is the most likely answer, then to assess whether or not the available evidence supports this answer. If it does, we pat ourselves on the back ("See, I knew it all along!") and don't look any further. The problem with this is that we often fail to recognize that most of our information is also consistent with other explanations or conclusions. We are easily led astray by the fact that there appears to be so much evidence to support our preconceived point of view.

ACH differs from conventional intuitive analysis in three important ways. Each of these differences is discussed in greater detail elsewhere in this Tutorial.

- ACH requires that you identify and analyze a full set of alternative hypotheses rather than a single most likely conclusion. This ensures that less likely but possible hypotheses receive fair treatment.
- You proceed by trying to refute or eliminate hypotheses, whereas conventional intuitive analysis generally seeks to confirm a favored hypothesis. The correct hypothesis is the one with the least—or no—inconsistent information.
- Instead of looking at one hypothesis and weighing the evidence pro and con for that hypothesis, you look at each item of evidence, one at a time, and assess whether that evidence is consistent or inconsistent with each of the hypotheses. This enables you to determine the "diagnosticity" of the evidence. An item of evidence is diagnostic when it helps you determine that one hypothesis is more likely to be true than another hypothesis. An item of evidence that is consistent with all hypotheses has no diagnostic value.

[Contents](#)

Refuting vs. Confirming Hypotheses

A fundamental precept of the scientific method is that one should proceed by rejecting or eliminating hypotheses, while tentatively accepting only those hypotheses that cannot be refuted. No matter how much information you have that is consistent with a given hypothesis, you cannot prove that hypothesis is true, because the same information may be and often is consistent with one or more other hypotheses. On the other hand, a single

item of evidence that is inconsistent with a hypothesis may be sufficient grounds for rejecting that hypothesis. A classic example is a criminal suspect who has a solid alibi. A natural human tendency is to give more weight to information that supports our favorite hypothesis than to information that weakens it. This is unfortunate, as we should do just the opposite.

The scientific method obviously cannot be applied *in toto* when working with ambiguous and incomplete information, but the principle of seeking to disprove hypotheses, rather than to confirm them, can and should be applied to the extent possible. Instead of being content with showing that your favored hypothesis is supported by a lot of evidence, you need to refute all the possible alternatives. Any hypothesis that you cannot refute should be taken seriously. This switch in perspective forces you to ask questions and seek evidence you would not otherwise consider. This is what provides insurance against unpleasant surprise.

The ACH approach forces you to spend more analytical time than you otherwise would on what seem to be the less likely hypotheses. This is desirable. Assessing the seemingly less likely hypotheses usually involves plowing new ground and, therefore, requires more work. What you start out thinking is the most likely hypothesis tends to be based on a continuation of your own past thinking. ACH helps to prevent surprise by forcing you to give a fairer shake to all the alternatives.

Because it is more difficult to refute hypotheses than to confirm a favored hypothesis, this approach will often increase your level of uncertainty. You may become less confident of what you thought you knew. Such uncertainty is frustrating but may be an accurate reflection of the true situation. The ACH procedure has the offsetting advantage of focusing attention on the few items of critical evidence that cause the uncertainty or which, if they were available, would alleviate it. This can guide future collection, research, and analysis to resolve the uncertainty and produce a more accurate judgment.

[Contents](#)

ACH Is a 9-Step Process¹

The full ACH process follows the following nine steps.

1. Define the question. Make sure that everyone involved in the process has a solid understanding of the specific question to be answered.
2. Identify all possible hypotheses, striving to assure that they are mutually exclusive. Use a group of analysts with different perspectives to brainstorm all possible hypotheses. Include the deception hypothesis when appropriate to ensure that you have exhausted all the possibilities.
3. Make a list of evidence and arguments for and against the hypotheses. Remember to

¹ The ACH methodology developed by Richards J. Heuer, Jr., consists of eight steps. We have added one more step at the outset to underscore the importance of defining the focal question.

include assumptions, logical deductions, and conclusions from other analyses – anything that affects your judgment about the likelihood of any hypothesis. Also include the absence of evidence one would expect to find if a hypothesis were true.

4. Prepare a matrix analyzing the “diagnosticity” of the evidence. Array the hypotheses across the top and the evidence down the side. Work horizontally to assess each data item’s consistency or inconsistency with each hypothesis. There is an option to also assess the Credibility and Relevance of each item of evidence to determine how much weight it should have in the analysis.
5. Draw tentative conclusions about the relative likelihood of each hypothesis, discounting evidence and arguments that have no diagnostic value. Try to refute hypotheses rather than confirm them. Does the Inconsistent evidence make a persuasive case for dismissing the hypothesis? Check the consistency of your evidence ratings. Analyze the sensitivity of your conclusions to a few items of highly diagnostic evidence. What are the consequences for the analysis if that item of evidence is wrong, misleading, or subject to a different interpretation?
6. Refine the matrix and reconsider the hypotheses. Identify gaps in the evidence that may need to be filled to refute hypotheses. Sort the evidence by type of evidence and be alert to any indication of possible deception. Solicit critical input from other analysts.
7. Compare your personal conclusions about the relative likelihood of the hypotheses with the Inconsistency Score or the Weighted Inconsistency Scores generated by the software. If they are not similar, figure out why and what you should learn from this. See guidance in the sections on [Interpreting the Inconsistency Scores](#) and [Potential Pitfalls](#).
8. Report conclusions, discussing the relative likelihood of all the hypotheses, not just the most likely one. Include discussion of why some alternative hypotheses were rejected or discounted.
9. Identify indicators or milestones for future observation to monitor whether one or more hypotheses might be changing.

[Contents](#)

ACH as a Collaborative Process

ACH is an excellent framework for collaboration among analysts. Although the software tool can be used by an individual analyst, the cross-fertilization of ideas when a group of analysts work together as a team helps analysts avoid personal bias and generates more and better ideas. When a team approach is adopted, the matrix can combine inputs and insights from analysts with different backgrounds. When analysts disagree, the matrix can be used to highlight the precise area of disagreement (for example, by using the Flag function). Subsequent discussion can then focus productively on the ultimate source of the differences. It is also possible to do a sensitivity analysis to see how alternative interpretations of the evidence or different assumptions affect the likelihood of the

alternative hypotheses. This often helps resolve, or at least narrow down, areas of disagreement.

When large numbers of analysts are involved, the process can be implemented with a four-step process:

- Identify an agreed-upon set of data relevant to the topic; remember to include assumptions, logical deductions, the absence of data, and conclusions from other analyses.
- Convene a structured brainstorming session with a diverse group of analysts to identify all the potential hypotheses.
- Commission a small group of analysts to work independently—and on their own time schedules—to load and analyze the data. This might take several days or weeks. In some cases, you might want to commission two small groups and have them work separately. This should provide independent validation of the key conclusions.
- Reconvene the larger group to assess the results of the working group. Focus on what data emerges as most diagnostic, the most persuasive reasons for discounting hypotheses, the credibility of the data supporting the most likely hypotheses, and the most productive areas for future research or collection.

[Contents](#)

When ACH May Not Meet Your Needs

ACH will stimulate and guide an inquiring mind but will not force open a closed mind. It assumes analysts are truly interested in identifying and questioning assumptions and in developing new insights about the issue in question. It is unlikely that ACH will resolve an impasse between analysts who are firmly entrenched in their views about the issue and have a strong commitment to maintaining those views. If an analyst is unable to see alternative perspectives, the evidence will always be interpreted in a way that supports that analyst's preferred view. ACH can still be useful, however, in helping to pin down the exact basis for the disagreement.

If your goal is mathematical accuracy in calculating probabilities for each hypothesis, there are other versions of ACH that may better meet your needs. They use Bayesian inference or Bayesian belief networks and may require a methodologist trained in Bayesian statistics to assist you through the process. Although the Bayesian probability calculations are mathematically correct, the results cannot be any more accurate than the multitude of subjective judgments about the evidence that go into the Bayesian calculation. This ACH software program emphasizes what is practical and understandable for the average analyst to use. Its payoff comes from the analytical process it takes the analyst through, not from precise probability calculations for each hypothesis. The final judgment is made by the analyst, not by the computer.

ACH is not appropriate for all types of decisions. It is used to analyze hypotheses about

what is true or what is likely to happen. If one wants to evaluate alternative courses of action, such as alternative business strategies, which computer to buy, or where to retire, this software is of limited value. The ACH matrix can be used to break such a decision problem down into its component parts, with alternative choices (comparable to hypotheses) across the top of the matrix and goals or values to be maximized by making the right choice (comparable to evidence) down the side. However, this type of analysis requires a different type of calculation. The principle of refuting hypotheses (in this case alternative courses of action) cannot be applied to a decision based on goals or personal preferences. One would need a more traditional analysis of the pros and cons of each alternative.

[Contents](#)

Project Management

Classification and Title

The classification of this document and the title of this analysis are entered in the space provided in the panel on the left side of the screen. When the analysis is printed, the left and right panels do not print. The printing is limited to the matrix plus the classification, title, and version name for the matrix which are printed above the matrix.

[Contents](#)

Adjust Column Width

If you need additional space to add more hypotheses, narrow the width of the columns or slide one or both of the panel dividers to make more room. Column widths within the matrix can be adjusted in the header (top cell) of each column. Click on the column line in the top cell and it can then be moved left or right. The Tutorial to the right of the matrix and the panel to the left of the matrix each have a sliding divider. Click anywhere on this divider and then move it left or right. You may also click on the Hide Tutorial button to remove the Tutorial from the screen.

[Contents](#)

Hide/Show Columns

To save space in the matrix, it is possible to hide or show certain columns that are used less frequently than others. This is done by clicking on the Hide/Show Columns button or selecting Hide/Show Columns in the Options menu and then checking the columns you want to appear in the matrix. The default value is to show all columns except the Date, Time, and Code columns.

[Contents](#)

Changing Colors and Font

Blue and yellow are the default colors for the evaluation of evidence under each hypothesis, but other colors are available. A red-green color scheme can be selected under the Options menu. Custom colors are also available by selecting Set Custom Colors in the Options menu. To change to a custom color, select Set Custom Colors, select the matrix element you wish to change, select the color from the color panel, click OK and then OK again.

A color change affects only the matrix in which the change is made. It does not change the blue and yellow default color.

[Contents](#)

Printing

To print a matrix, select Print in the File menu. If you have created multiple matrices, the program will only print one matrix at a time. To print all matrices, they must be selected and printed individually. You may need to adjust the column widths or slide the panel dividers so that the printed matrix will fit on one page.

The Tutorial is printed separately. To print it, click at the top of the Tutorial where it says Print Tutorial. This shows a Print Preview that enables you to print either the full Tutorial or only selected pages. The menu at the top of the Preview shows the page numbers. You can also select whether to Print the Tutorial in Portrait or Landscape format.

To print Evidence Notes, there are two options that can be selected in the File menu. They are Print Evidence and Print Preview of Evidence. Print Evidence prints the notes for all items of evidence. Print Preview of Evidence gives you a visual preview of the print out so you can select which pages you want to print. These functions print both the short version of the evidence and the Evidence Notes.

[Contents](#)

Building the Matrix

Enter Hypotheses

Deciding what hypotheses to evaluate is the first step in ACH. Hypotheses must be carefully considered and worded, as they form the foundation on which the analysis is built. Input from several different analysts with different perspectives is strongly encouraged.

A hypothesis is a testable proposition about what is true, or about what has, is, or will happen. It should usually be worded as a positive statement that can be disproved. A

good set of hypotheses meets two tests: strive to make the list of hypotheses both **comprehensive** and **mutually exclusive**. The hypotheses should cover all reasonable possibilities, including those that seem unlikely but not impossible. If one hypothesis is true, then all other hypotheses should be false. Sometimes it is useful to think of a hypothesis as a set of unique stories of how an event played out or will play out. As evidence is collected and added to the matrix, you may find that hypotheses need to be reworded, refined, split into two hypotheses, combined, added, or deleted.

When deciding whether to include an unlikely hypothesis, consider whether the hypothesis is virtually impossible or simply unproven because there is no evidence for it. For example, the possibility that an adversary is trying to deceive you should not be rejected just because you see no evidence of deception. If deception is done well, you should not expect to find evidence of it readily at hand. The possibility should not be rejected until it is disproved, or, at least, until after you have made a systematic search for evidence and found it lacking.

Each new project starts with columns for two hypotheses. To enter text for the first hypothesis, do not click on the Enter Hypothesis button, as this will simply create an additional hypothesis column. Instead, click on the empty cell below the H: 1. Type in a brief description of the hypothesis, then hit Enter on your keyboard. If you forget to hit Enter, the hypothesis will disappear and you will have to type it in again. Then do the same in the cell below H: 2. To create additional hypotheses, click on the Enter Hypothesis button. To edit a hypothesis, click once on the hypothesis description and then make the changes. To delete a hypothesis, click once on the hypothesis number above the description to highlight the column, then hit the delete key. Alternatively, place the cursor anywhere in the column and select Delete Selected Hypothesis in the Matrix menu.

When evidence is entered and then sorted by diagnosticity or any other characteristic, the hypotheses are reordered by Inconsistency Score or Weighted Inconsistency Score, with the most likely hypothesis to the left. To facilitate consistent entering of data, a click on the Enter Evidence button returns both the evidence and the hypotheses to their original order.

[Contents](#)

Enter Evidence

The word evidence is interpreted very broadly. It refers to all factors that influence your judgment about the hypotheses. This includes assumptions and logical deductions as well as specific items of intelligence reporting. Assumptions or logical deductions about capabilities, intentions, or how things normally work in the foreign country or culture involved are often more important than hard evidence in determining analytical conclusions.

The absence of evidence is also evidence and must be noted. For example, if you are

analyzing the intention of an adversary to launch a military attack, the steps the adversary has **not** taken may be more significant than the observable steps that have been taken. Ask yourself the following question for each hypothesis: If this hypothesis is true, what are all the things that must have happened, or may still be happening, and what evidence of this should I expect to see? Then ask: If I am not seeing this evidence, why not? Is it because it is not happening, it is being concealed, or because I have not looked for it?

Not all evidence needs to be included. For some types of issues, such as warning analysis, it may be useful to prune the older evidence. A collection of older evidence is likely to bias the analysis in favor of concluding that the status quo will continue. If there is going to be a significant change (such as a military attack, a coup d'etat in the near future, or a surprising election victory or defeat) that may well be apparent only from the recent evidence.

If you are uncertain whether an item of evidence is true or deceptive, it is often advisable to enter that evidence twice, once with the assumption that it is not deceptive, and once with the assumption that it is deceptive. For example, a foreign leader makes a public statement of intentions, such as "we have no interest in developing weapons of mass destruction." You cannot rate the consistency or inconsistency of that statement with your hypotheses without pre-judging whether the foreign leader is being truthful or trying to hide a weapons development program. The solution is to enter and rate that item of evidence twice, once with the assumption that it is true, and once with the assumption it is deceptive. Entering this evidence twice with different consistency ratings forces you to think about the deception, and this is preferable to entering it only once and rating it as consistent with all hypotheses. In this case, you may also want to put a check in the flag column for both entries to remind you that you have entered the same evidence item twice.

To enter evidence, click on the Enter Evidence button and put the cursor in the evidence cell. If some evidence has already been entered and sorted for analysis, a click on the Enter Evidence button automatically returns the evidence to the order in which the evidence was entered, returns the hypotheses to their original order, and hides all inconsistency score calculations so that they do not influence your judgments about the consistency of the evidence.

For analysts who maintain lists of evidence in an Excel database, it is possible to cut and paste a whole column of evidence instead of entering one item of evidence at a time. To paste a column of evidence from Excel, first copy the column of evidence in Excel, select the Evidence column in ACH, and paste with CTRL-V or by selecting Paste in the Edit menu.

To edit evidence, click once on the evidence description and then make the changes. To see or edit the Notes, click twice on the evidence description. To delete an item of evidence, click once on the number at the far left of the row to highlight the row, then delete. Alternatively, place the cursor anywhere in the row and select Delete Evidence in the Matrix menu.

Each item of evidence is automatically numbered as it is entered. This reflects the order in which evidence was entered, and this numbering cannot be changed. The evidence can be sorted in various ways (see [Sort Evidence](#)), and this changes the order of the evidence. You can also return the evidence to the original order by going to Sort Evidence and selecting Order Added.

Because the order of evidence is changed whenever it is sorted, it is sometimes difficult to locate a specific item of evidence. In this case, go to the File menu and select Find Evidence. This searches each piece of evidence sequentially. It first shows matches in the Evidence name column. Then it shows matches in the Evidence Notes panel.

[Contents](#)

Evidence Notes

The Evidence Notes panel to the left of the matrix is for entering information beyond what can be put in the matrix. The Notes may include a fuller description of the evidence, a reference to the source of the evidence, a hyperlink directly to the source document, or extracts or a full copy of the report. An alternative or additional use for Evidence Notes is to comment on why a particular rating was given for a particular hypothesis. You can provide text explaining the rating in each cell. You can also flag differences of opinion relating to why a specific rating was given to a particular piece of evidence.

Clicking on any item of evidence provides a place for entering Notes specifically on that item. Once a Note has been entered, that Note is visible whenever the cursor is in that row of evidence. There is no requirement that each item of evidence have a Note. The presence of a Note in any item of evidence is indicated by a paper clip symbol immediately to the right of the evidence number.

You can enter in the Evidence Notes a hyperlink or hyperlinks to source documents for this item of evidence. The source document comes up in a new window. You can link to any document that you can reach with your browser. You can either manually type in the link or cut and paste the link text from your browser. Note that links must start with one of `http://`, `https://`, `ftp://`, or `file:///`. To link to a file on your local computer, your full file path should start with something like `file:///`, followed by your hard drive designation and other identifiers. Note that spaces in the file path are not supported. Instead, you can enter %20 for spaces (e.g. `file:///C:/Program%20Files/Windows%20NT/Pinball/PINBALL.EXE`) or simply use the DOS 8.3 version of the path (e.g. `"C:/Docume~1/"` for `"C:/Documents and Settings/"`). To follow a link, either double click or press CTRL and click at the same time. Dragging and dropping links is not supported.

Two options for printing Evidence Notes can be selected in the File menu. They are Print Evidence and Print Preview of Evidence. Print Evidence prints the Notes for all items of evidence in one large file. Print Preview of Evidence gives you a visual preview of the

print out so you can select which Notes you want to print. The menu at the top of the Preview shows the page numbers.

[Contents](#)

Enter Type of Evidence

Type of evidence usually refers to the type of source that provided the evidence, but any other analytically useful set of categories may be used. Simply enter the category in the appropriate column of the matrix. To sort by type of evidence, go to Sort Evidence and select Type. This will sort the types of evidence alphabetically.

For intelligence analysis, common types of evidence include such categories as HUMINT, SIGINT, Imagery, Open Source, Assumption, Logical Deduction, and Absence of Evidence. In a criminal investigation, it may be appropriate to have categories such as Police Report, Eye Witness Account, Interview, and Forensic Evidence. A counterintelligence investigator might use the column to catalogue Motive, Opportunity, Means, and Character Assessment.

When the evidence is sorted by Type, the evidence listed within each category is also sorted by diagnosticity. This allows you to identify, for example, the most diagnostic or discriminating items of HUMINT or Imagery. Sorting and analyzing the evidence by type can provide clues to the reliability of sources and signal possible deception. If all types of sources are telling a consistent story, that is a good sign. If not, try to figure out why. Are some sources vulnerable to manipulation for the purpose of deceiving you?

[Contents](#)

Enter Date or Time

The date or time when events occurred is significant for some analyses, especially counterintelligence or criminal investigations. The Date and Time columns are the place to enter this kind of information. The default status for these functions is to hide them in order to simplify the initial presentation of the matrix and make more room to list hypotheses. To show the Date and/or Time column so that you can enter data, go to the Hide/Show Columns button or select Hide/Show Columns in the Matrix menu. Then check the box for Date or Time, or both, as desired. To sort by Date or Time, go to Sort Evidence and select the desired column. Sorting by Date or Time converts all your evidence into a chronology.

Setting a Date: A single click in a Date cell enters the current day's date and provides two different methods for changing the date -- a button and a scroller. A single click on the button brings up a calendar for setting day, month, and year. Clicking on the up or down arrow on the scroller also changes the day, month, or year, depending on where the cursor is placed. To change the day, for example, place the cursor in the day part of the date. To

delete a date, highlight and delete the date, then remove the scroller by moving the cursor to any other cell outside the date column.

Setting the Time: A single click in a Time cell enters the time 12:00 pm and presents buttons for scrolling earlier or later. The default setting is to scroll by the minute rather than the hour, so set the minutes first. If you then put the cursor in the hour area, it will scroll by the hour. To delete a time, highlight and delete the time, then remove the scroller by moving the cursor to any other cell outside the date column.

[Contents](#)

Enter Code

This is an extra column in the matrix. The Code function gives the analyst the flexibility to enter any additional set of categories the analyst might wish to use to sort the evidence. For example, you might want to code the data by country or by state. You can also use this column to track a specific stream of HUMINT reporting, or to identify reporting from a specific news organization like the BBC. The Code function can also be used to flag specific reports that you suspect might constitute denial, deception, or inaccurate reporting. The Code column sorts alphabetically or numerically.

The default status for this function is that the matrix column is hidden. To show the Code column, go to the Select Column button or the Select Column option in the File menu. To sort by Code, go to Sort Evidence and select Code.

[Contents](#)

Assess Consistency of Evidence

After entering an item of evidence, the next step is to assess whether that item of evidence is Consistent or Inconsistent with each hypothesis. When an item of evidence is first entered, the default value for the Consistency rating for that item of evidence is set to NA for Not Applicable. For each hypothesis, ask yourself: if this hypothesis were true, is it likely that I would see this evidence? If the answer is "Yes," change the consistency rating to show that the evidence is Consistent (C) with the hypothesis; if "No," mark it as Inconsistent (I). If it is Very Consistent or Very Inconsistent, mark it CC or II. The Very Consistent (CC) and Very Inconsistent (II) ratings are used when you have a high degree of confidence in providing the rating or believe the item of evidence makes a compelling case supporting or contradicting a hypothesis.

Evidence may also be Neutral (N) or Not Applicable (NA) to some hypotheses. The Neutral (N) designation is used when there are alternative interpretations of the evidence, one Consistent and the other Inconsistent. The Not Applicable (NA) rating is appropriate when an item of evidence is clearly not applicable to one or more of the hypotheses.

To change a Consistency rating, single click on the appropriate cell and then select the desired setting from the menu of options.

Working across the matrix. In entering the Consistency ratings, it is essential that you work *across* the matrix, assessing the Consistency of the evidence, one item at a time, against each of the hypotheses. Do not work *down* the matrix. That is, do not take one hypothesis at a time and assess the consistency or inconsistency of all the evidence for that single hypothesis.

This procedure enables you to assess what is called the "diagnostic" value of the evidence. Diagnosticity of evidence is an important concept that is, unfortunately, unfamiliar to many analysts. Evidence is diagnostic when it is Inconsistent with one or more hypotheses and Consistent with others. That is, it influences your judgment on the *relative* likelihood of the various hypotheses. If an item of evidence is Consistent or Inconsistent with *all* hypotheses, it has no diagnostic value. In doing an ACH analysis, it is a common experience to discover that much of the evidence supporting what you believe to be the most likely hypothesis is really not helpful, because the same evidence is consistent with all the other hypotheses. In this case, it would be misleading to use this particular item of evidence to support your analytic conclusion. In some cases, it might even prove counterproductive because someone reading your assessment who supports a contrary hypothesis could argue that the evidence you cite also is consistent with their view.

If an item of evidence is Very Inconsistent (II) with a hypothesis, and if that evidence rates high on Credibility and Relevance, this is a strong indicator that the hypothesis is unlikely. To sort by Diagnosticity, go to Sort Evidence and select Diagnosticity.

The standards you use for judging Consistency, Credibility, and Relevance sometimes evolve as you gain a better understanding of the relationship between the hypotheses and the evidence. After entering all the evidence, go back and make sure the judgments are consistent. Change any ratings you now see in a different light. You may need to do this several times during the course of the analysis.

ACH uses the Consistency ratings to calculate an Inconsistency Score for each hypothesis. Evidence that is Inconsistent (I) counts one point, and Very Inconsistent (II) counts two points. Because the focus is on refuting hypotheses rather than confirming them, only Inconsistent and Very Inconsistent evidence is counted in the Inconsistency Score. Although the software is not sensitive to whether an item of evidence is rated Consistent (C) or Very Consistent (CC), you can still factor this into your analysis if it helps you discriminate between key items of evidence. There is also a Weighted Inconsistency Score that takes the Weight of evidence into account as described below.

[Contents](#)

Assess Weight of Evidence

The software enables the analyst to assign weights to each item of evidence. This is optional because it adds an element of complexity to the analysis that is not always needed to achieve the basic goals of the ACH analysis. How these weights are conceptualized is also optional. The matrix has columns for two weight variables. These are labeled Credibility and Relevance. Definitions and examples of these concepts are provided below. However, these are somewhat murky concepts, alternative ways of defining weight variables have been suggested, and the best choice may depend upon the subject matter and type of evidence available. For example, experience suggests that use of the Credibility and Relevance functions are more likely to prove useful in criminal investigations than in counterintelligence investigations. The most important consideration is whether the analyst feels comfortable using these functions and believes they enhance the quality of the analysis.

Although the titles of the columns cannot be changed, you may if desired redefine these weight variables if that better suits your needs. The only constraint is that their use be consistent with how these variables contribute to the calculation of the Weighted Inconsistency Scores as described in [Calculation of Scores](#). Briefly, each of the two weight variables can be rated as High, Medium, or Low for each item of evidence. A High weight for both variables, e.g., Credibility and Relevance, doubles the value of an Inconsistent or Very Inconsistent rating of the evidence, while a low weight for both cuts the value in half. When Credibility and Relevance or their equivalents have different ratings, an intermediate value is used. This means that the value attributed to a single item of evidence in the Weighted Inconsistency Score can range from 0.5 to 4. For further information on the calculation, see [Calculation of Scores](#).

If you are uncomfortable using the Credibility or Relevance weights, you can decide to start the analysis without them and then add the weights later if desired. Another option is to collapse the Credibility and Relevance variables into a single Weight variable for each item of evidence. You do this simply by entering the same weight, High, Medium, or Low, in both the Credibility and Relevance columns. In other words, you do not need to distinguish between Credibility and Relevance if you don't think that is necessary or appropriate for your analysis.

The default setting for a new project is that both the Credibility and Relevance columns are showing in the matrix. If you do not want to use these columns, go to Type of Calculation and change it to Inconsistency Score. This type of calculation does not use the weight variables, and those columns will disappear from the matrix. Alternatively, you can go to the Hide/Show Columns button to either hide or show either one or both of these columns.

[Contents](#)

Credibility of Evidence

This column is shown only when the Type of Calculation is set for Weighted Inconsistency Score. If you want to assess the Credibility of each item of evidence, use

this column to assign a weight of High, Medium, or Low as discussed under [Assess Weight of Evidence](#).

Credibility depends on both the reliability of the source and the credibility of the information provided by the source, which are separate but closely related concepts. For example, the reliability of information posted on the Internet varies widely as we all know. Even a highly reliable source, however, can produce information of questionable accuracy. A proven source may report inaccurate information if the source has only indirect or incomplete access to the information, the source is accurately reporting information derived from a subsource of unknown reliability, or if the source is reporting his or her own conclusions which may or may not be accurate. Similarly, imagery and SIGINT are unimpeachable sources but may be vulnerable to denial or deception if our collection capability is known to the other side. Some technical sources are also vulnerable to faulty interpretation of imagery or faulty translation of communication intercepts. For an eyewitness source in a criminal investigation, credibility of the reported observation depends upon such things as time elapsed since the observation, circumstances of how the observation was made, and whether the witness at the time of the observation recognized the importance of remembering what was observed. The analyst should consider such factors before rating Credibility.

Sorting by Credibility highlights the evidence that you have said is most credible. It is an alternative means of identifying evidence that may warrant a hard look at how the evidence was evaluated. To sort by Credibility, go to Sort Evidence and select Credibility. For a discussion of how this variable affects the Weighted Inconsistency Score, see [Calculation of Scores](#). If you do not want to assign Credibility or Relevance weights, change the scoring calculation to Inconsistency Score and this column will disappear.

[Contents](#)

Relevance of Evidence

This column is shown only when the Type of Calculation is set for Weighted Inconsistency Score. If you want to assess the Relevance of each item of evidence, use this column to assign a weight of High, Medium, or Low as discussed under [Assess Weight of Evidence](#).

One assesses the Relevance of the evidence to the question being analyzed, not the relevance to each individual hypothesis. The analyst asks whether this evidence is central to the analytical judgment that must be made or only indirectly related. Consider the example of Country X, which has moved its troops to the border of Country Y. The hypotheses are that Country X plans to attack unless it gets its way in a dispute with Country Y, or that Country X has decided not attack even if Country Y does not give in to Country X's demands. A report from a source on Country X's General Staff saying that the attack will start next week is obviously Very Inconsistent with the No Attack hypothesis. On the other hand, a report that one regiment of Country X's army stationed

along the border has cancelled all leaves is also Inconsistent with the No Attack hypothesis and may be from a very credible source, but it is less directly related to the issue of Attack or No Attack. Therefore, it would rank lower in Relevance. Comparable judgments are made by counterintelligence and law enforcement analysts. Evidence that points to a specific individual as the perpetrator is highly Relevant. Circumstantial evidence about motive, opportunity, or means may be equally Consistent or Inconsistent with certain hypotheses, but it may rate lower on Relevance because its evidential value is less decisive.

Sorting by Relevance highlights the evidence that you have said is most relevant. It is an alternative means of identifying evidence that may warrant a hard look at how the evidence was evaluated. To sort by Relevance, go to Sort Evidence and select Relevance. For a discussion of how this variable affects the Weighted Inconsistency Score, see [Calculation of Scores](#). If you do not want to assign Credibility or Relevance weights, change the scoring calculation to Inconsistency Score and this column will disappear.

[Contents](#)

Flag Column

The Flag column is used to mark those items of evidence where there are questions or disagreement about the most appropriate weight or consistency rating for that evidence. Discussion of alternative evaluations of the evidence is encouraged, as it gives analysts the opportunity to learn about alternative interpretations and the rationale for them. When there is disagreement over how to rate a particular item of evidence, it is often more efficient to flag the item and move on to score the remaining items of evidence. Once the entire matrix is completed, the analysts then can go back and revisit flagged items. In some cases, subsequent discussions clarify the issue, making it easier to score the flagged item. In other instances, the item of evidence turns out to have little diagnostic value suggesting that there is little merit arguing the case. If, on the other hand, the item has significant diagnostic value, then the debate should be vigorously pursued with opposing arguments rigorously laid out by each side.

The Flag column is located to the right of the hypotheses. Clicking once in the empty cell adds a red flag. A second click changes it to a blue flag, and the third click deletes the flag. Sorting by flags groups that evidence together to facilitate further analysis or discussion. The red flag might be used to indicate evidence that needs to be discussed while the blue flag could indicate it has been discussed.

[Contents](#)

Analytical Functions

Sort Evidence

The sorting of evidence to highlight or compare certain kinds of evidence is a principal analytic function. The Sort Evidence By button allows the analyst to sort by Diagnosticity, Weight, Type, Date or Time, Flags, Code, and Order Added. Each sort, except for Order Added which returns the evidence to its original order, recalculates the hypothesis scores as necessary and reorders the hypotheses with the most likely hypotheses to the left. Each of these sorts is described briefly below.

Diagnosticity: Sorting by Diagnosticity brings to the top of the matrix the evidence that, according to your ratings, has the greatest impact on the relative likelihood of the hypotheses. This is the Inconsistent and Very Inconsistent evidence, which is the only evidence used in the ACH scoring. By looking carefully at this group of evidence, you should gain a good understanding of what evidence is driving the Inconsistency Scores. This gives you an opportunity to reaffirm or reconsider your evaluation of the consistency and the weight of this evidence. As you change consistency ratings or weights, observe how much difference it makes in the score of each hypothesis. If the Inconsistency Score or Weighted Inconsistency Score does not reflect your thinking about the most and least likely hypotheses, see below on Interpreting the Inconsistency Scores.

When sorting by Diagnosticity, the software assigns each row of evidence a weight determined by the number of Inconsistent and Very Inconsistent cells on that particular row. The rows with the highest scores appear at the top of the matrix and those with the lowest scores at the bottom. If several rows have the same score, the rows are listed in order entered. This ranking by diagnosticity does not take into consideration any differences in how a particular item of evidence was rated for Credibility or Relevance.

Credibility and Relevance: Evidence that is Credible and Relevant is potentially significant in making your analytical judgments. Sorting on these variables is an alternative means of identifying evidence that may warrant a hard second look at how the evidence was evaluated.

Type: Type usually refers to type of source. Sorting by Type of source can provide clues to the reliability of sources and possible deception. If all types of sources are telling a consistent story, that is a good sign. If not, try to figure out why. Are some sources vulnerable to manipulation for the purpose of deceiving you? The types of evidence are sorted alphabetically. Within each type of evidence, the evidence is sorted by diagnosticity. This, for example, allows you to assess independently which HUMINT or which Open Source evidence is the most diagnostic.

Date and Time: Sorting by Date and/or by Time enables analysts to examine the sequence in which events occurred. This is significant for some analyses, especially counterintelligence or criminal investigations. When evidence is entered, it is often not entered in chronological order. The Date and Time functions allow analysts to resort their data and generate a true chronology of events.

Flags: Flags are used to identify evidence where there are questions or disagreement

about the most appropriate weight or consistency rating for that evidence. Sorting by Flags groups that evidence together to facilitate further analysis or discussion.

Code: The Code function gives the analyst the flexibility to enter any *additional* set of categories the analyst might wish to use to sort the evidence. Entries in this column are sorted alphabetically or numerically.

Order Added: This is the numerical order in which evidence was entered. Any other sorting changes this order. Whenever you click on Enter Evidence, this returns evidence to the original order to facilitate entering of new evidence. You can also return the evidence to the original order by going to Sort Evidence and selecting Order Added.

[Contents](#)

Create Duplicate Matrix

The Duplicate Matrix button creates a duplicate matrix that you can use for various tests or experiments without changing the original matrix. Test how a series of changes in the evaluation of evidence or the weights affect the scores of the hypotheses. Two analysts using the same evidence could evaluate that evidence differently, in separate matrices, and then see how much difference this makes in the scores of the hypotheses.

For convenience in shifting from one matrix to another, multiple matrices are kept in a single project. The panel to the left of the matrix has a list of all available matrices. To navigate from one matrix to another, simply click on the name of the desired matrix. If you have multiple matrices and want to create another, you will need to designate which matrix should be duplicated. To do this, simply go to the matrix you want to duplicate before clicking the Duplicate Matrix button.

The original matrix is named Main, and subsequent matrices are named Test 1, Test 2, etc. To change the name of any matrix, double click on the name and then enter the new name in the dialog box. When a matrix is printed, the version name will be printed above the matrix next to the title. You may only print one matrix at a time. It is not possible to print the entire project.

To delete a matrix, highlight the matrix name; then click on your keyboard Delete button and click on Yes to confirm the delete. Alternatively, highlight the matrix name and then select Delete Selected Matrix in the Matrix menu.

[Contents](#)

Partition Evidence

Partitioning the evidence refers to dividing the evidence into two groups that may be analyzed separately and compared. For example, you can compare recent evidence

against older evidence. If the Inconsistency Scores differ for these two groups of evidence, it suggests that the situation may be changing. You may decide to delete older evidence. To assess the possibility of deception or the reliability of sources, you can compare open source reporting versus clandestine source reporting or clandestine human source reporting versus SIGINT reporting.

It is often advisable to compare the Inconsistency Scores based only on your assumptions and logical deductions with the Inconsistency Scores based only on the hard evidence. For example, it is useful for Indications and Warning analysis to compare strategic assumptions about the adversary's capabilities and intent with tactical evidence of what is actually happening in the field. If one group of evidence points to a different conclusion than the other, consider why and evaluate the significance of this disparity. If the most diagnostic evidence consists largely of assumptions and deductions rather than hard evidence, confidence in your conclusions should be limited.

To partition the evidence, click on "Duplicate Matrix" to create a new, identical matrix. Then delete evidence as appropriate; for example, delete the "old" evidence from one and the "recent" evidence from the other. When making such a partition based on the timeliness of the evidence, evidence such as assumptions and logical deductions that apply to both time periods must be included in both partitions.

[Contents](#)

Calculation of Scores

This software makes two different calculations regarding the likelihood of each hypothesis. The Inconsistency Score is a simple count of the number of items of Inconsistent (I) and Very Inconsistent (II) evidence for each hypothesis. The Weighted Inconsistency Score follows the same principle but adjusts the score according to the Credibility and Relevance of the evidence.

Inconsistency Score: Consistent with the principle of refuting hypotheses, only the Inconsistent evidence is counted. An Inconsistent (I) item of evidence counts as 1 and a Very Inconsistent (II) item counts as 2. The Inconsistency Score for each hypothesis is the sum of the Inconsistent (I) and Very Inconsistent (II) ratings for that hypothesis. The hypothesis with the highest Inconsistency Score is the least likely. The one with the lowest Inconsistency Score is the most likely.

Weighted Inconsistency Score: This is like the Inconsistency Score, except that the value of each Inconsistent (I) or Very Inconsistent (II) rating is adjusted according to the Credibility and Relevance of that item of evidence. When you tell the software to calculate a Weighted Inconsistency Score, it selects a factor based on the Credibility and Relevance assigned to that line of evidence. That factor then is multiplied against the value for the selected Inconsistent (I) or Very Inconsistent (II) evaluation with a value of 1 assigned to Inconsistent (I) and a value of 2 assigned to Very Inconsistent (II). The factors for Credibility and Relevance are shown in the following matrix.

	Low Credibility	Medium Credibility	High Credibility
Low Relevance	(LL) $0.5 \times 0.5 = \mathbf{0.25}$	(ML) $0.7 \times 0.7 = 0.5$	(HL) $1.0 \times 1.0 = \mathbf{1.0}$
Medium Relevance	(LM) $0.7 \times 0.7 = 0.5$	(MM) $1.0 \times 1.0 = \mathbf{1.0}$	(HM) $1.4 \times 1.4 = 2.0$
High Relevance	(LH) $1.0 \times 1.0 = \mathbf{1.0}$	(MH) $1.4 \times 1.4 = 2.0$	(HH) $2.0 \times 2.0 = \mathbf{4.0}$

If the Credibility and Relevance are both rated Medium, the value of the Inconsistent remains at 1 and the value of Very Inconsistent (II) remains at a 2. If both the Credibility and Relevance are High, the value of the consistency rating is doubled; Inconsistent (I) then has a value of 2, and Very Inconsistent (II) has a value of 4. If both Credibility and Relevance are Low, the value of an Inconsistent or Very Inconsistent rating is cut in half to .5 and 1, respectively. If the Credibility and Relevance ratings differ, for example if Credibility is High but Relevance is Medium, a mathematical formula is used to calculate the values for all the possible combinations of Low, Medium, and High weights. The value of 0.7 (actually 0.71717) for ML and LM scores and 1.4 (actually 1.41414) for MH and HM scores are used for the intermediate values in order to ensure consistent ratings when you multiply the Credibility Score by the Relevance Score. We recommend against placing undue emphasis on the Inconsistency Scores. If you must use them, then go out only one decimal place. But recall the following caution: The scores give you a rough approximation of the most to least likely hypotheses. You—not the computer—must make the final decision as to which hypotheses can be discounted and which deserve priority attention.

The score for a single item of Inconsistent (I) evidence can range from a value of .5 for an Inconsistent (I) rating with Low Credibility and Low Relevance to a value of 4 for a Very Inconsistent (II) rating with High Credibility and High Relevance. This means that a highly credible item of Very Inconsistent (II) evidence that gets to the heart of the analytical problem may count eight times as much as an Inconsistent item of uncertain Credibility and peripheral Relevance.

Whenever an Inconsistent rating or a Credibility or Relevance weight is changed, the Inconsistency Score or Weighted Inconsistency Score is automatically recalculated. If this recalculation changes the likelihood ranking of the hypotheses, the order of the hypotheses is automatically reordered from the most likely (on the left) to the least likely on the right. If the evidence needs to be resorted, the words "Resort Needed" will appear in the Sort Evidence By window. If only the Credibility and Relevance weights are changed (and the Inconsistent scores remain unchanged), the order of the rows may not change. Diagnosticity sorts only by the weights determined by the number of Inconsistent (I) and Very Inconsistent (II) cells in the row.

[Contents](#)

Show Graphics

To show a graphic display of the Inconsistency or Weighted Inconsistency Scores for the hypotheses, select Show Chart on the Matrix menu. To close the graphic, click on the X

in the upper right corner.

[Contents](#)

Drawing Conclusions

Interpreting the Inconsistency Scores

It would be a serious mistake to think of the Inconsistency and Weighted Inconsistency Scores as the product of a magic formula that tells you which hypothesis is most likely. The Inconsistency Scores only tell you what you told the program to say through your selection of hypotheses and evidence and your evaluation of that evidence. The scores give you a rough approximation of the most to least likely hypotheses, but you—not the computer—make the final decision on which hypotheses can be discounted and which deserve priority attention.

One purpose of the Inconsistency and Weighted Inconsistency Scores is to focus your attention on the need to refute hypotheses, rather than to confirm them. The Inconsistency scores help you track your progress in accomplishing that task.

Another reason for calculating Inconsistency Scores is to use the scores as an independent check on your own thinking. After completing the matrix, compare the Inconsistency or Weighted Inconsistency Scores with your personal views about the hypotheses that you have formed as a result of all your research. If the relative likelihood of the hypotheses as shown by the scores matches up reasonably well with your own conclusions, this indicates that the matrix is an effective representation of your thinking. That is good. If there is a significant discrepancy between the computer-generated scores and your own thinking about the hypotheses, you need to figure out why. There are three principal reasons why your thinking may diverge significantly from the Inconsistency or Weighted Inconsistency Scores.

- If the scores in the matrix do not support what you believe is the most likely hypothesis, your matrix may be incomplete. Your thinking may be influenced by assumptions or logical deductions that have not been included in the list of evidence. If so, include them so that the matrix fully reflects everything that influences your judgment on this issue. It is important that you recognize the role that unstated or unquestioned (and sometimes unrecognized) assumptions play in your analysis. In political or military analysis, for example, your conclusions will often be driven by assumptions about another country's capabilities or intentions. A principal goal of the ACH process is to identify those factors that drive your thinking on the issue, so that they can then be questioned and, if possible, validated.
- If you think the scoring gives too much credibility to one or more of the less likely hypotheses, it is because you have not assembled the Inconsistent evidence needed to refute them. This may be because you devoted insufficient attention to obtaining such evidence, or because the evidence simply is not there. If you cannot find such

evidence, you may need to adjust your thinking to recognize that the uncertainty is greater than you thought.

- There are several circumstances described below under Potential Pitfalls when the Inconsistency and Weighted Inconsistency Scores can be misleading.

[Contents](#)

Potential Pitfalls

There are several circumstances when simply counting the amount of Inconsistent evidence can be misleading. A sound understanding of the ACH process and of your evidence should enable you to recognize and adjust to these circumstances. For example:

- *Definitive Evidence:* There are occasions when intelligence collectors obtain information from a trusted and well-placed inside source. The ACH analysis can assign High weights for Credibility and Relevance, but this may not be enough to reflect the conclusiveness of such evidence and the impact it should have on your thinking about the hypotheses. In other words, there may be circumstances when one or two authoritative reports clearly support one hypothesis and are sufficient to refute all other hypotheses regardless of what other less reliable or less definitive evidence may show. A classic example of this is a perfect alibi in a criminal murder investigation, when one "Inconsistent" data point is sufficient to refute all the other evidence supporting a hypothesis.
- *Unbalanced Set of Evidence:* Evidence needs to be representative of the problem as a whole. If there is considerable evidence on a related but peripheral issue and comparatively few items of evidence on the core issue, the Inconsistency or Weighted Inconsistency Scores may be misleading.
- *Diminishing Returns:* As evidence accumulates, each new item of Inconsistent or Very Inconsistent evidence has less impact on the Inconsistency Scores than the earlier evidence. For example, the impact of any single item is less when there are 50 items than when there are only 10 items. To understand this, consider what happens when you calculate the average of 50 numbers. Each number has equal weight, but adding a 51st number will have less impact on the average than if you start with only 10 numbers and add one more. Stated differently, the accumulation of evidence over time slows down the rate at which the Inconsistency Scores change in response to new evidence. Therefore, these numbers may not reflect the actual amount of change in the situation you are analyzing. If you are analyzing change over time, it is desirable to delete periodically the older evidence or to partition the evidence and analyze the older and newer evidence separately.

[Contents](#)

Why No Probability Calculation?

It may seem to some users that this ACH software could and should convert the rating of evidence in the matrix into numerical probabilities for each hypothesis rather than a simple Inconsistency or Weighted Inconsistency Score. That could be done, but it is deliberately *not* done for several reasons.

- A much more complex rating system using Bayesian statistical inference would be required to do it accurately, and the conclusions would still be no more accurate than the analyst's subjective judgments that go into rating the evidence. Other versions of ACH with greater mathematical precision have been developed for specialized purposes, but they are less practical for broad use and not easily understood by the average intelligence analyst. The two greatest drawbacks are that they require much more time to rate the evidence and, in most cases, the assistance of an expert facilitator.
- It would be too easy to misuse any simple probability calculation. This software is best used as an analytical aid, not a magic answer machine. Its purpose is to help analysts structure complex problems in ways that lead them to ask questions they might not otherwise ask. Learning and new ideas occur as analysts identify alternative hypotheses, examine the relationship between the evidence and each of the hypotheses, seek out evidence to refute hypotheses, and as they compare and discuss these matters with other knowledgeable analysts. The ultimate conclusion must come from the analyst, not a software algorithm.

[Contents](#)

Sensitivity Analysis

To analyze how sensitive your conclusion is to just a few critical items of evidence, sort the evidence by diagnosticity. Identify those items that have the greatest impact on your thinking and on the Inconsistency Score calculations. Check how much the Inconsistency Scores would change if that evidence were wrong or misleading or subject to a different interpretation.

For those few items of evidence that drive the outcome of the analysis in one direction or the other, go back and recheck the original evidence. Are there alternative explanations or interpretations? Misunderstanding or error in translation due to language problems is not uncommon. How reliable is the source? Could the evidence have been fabricated in a deliberate effort to influence your thinking?

When intelligence analysis turns out to be wrong, it is often the case that key evidence or key assumptions went unchallenged and was later proven invalid. It is a truism that analysts should identify and question their assumptions, but this is easier said than done. A major benefit of ACH is that it identifies the few items of decisive evidence, including assumptions and logical deductions, which need to be subjected to the most rigorous scrutiny.

[Contents](#)

Identify Milestones

Analytical judgments based on incomplete and ambiguous information should always be regarded as tentative. The situation may change, or it may remain unchanged while you receive new information that changes your appraisal of the situation.

Analysts should use the results of the ACH analysis to identify milestones to watch for which, if observed, would suggest a significant change in the Inconsistency Scores. Put these out as intelligence requirements to the collectors who are tracking the situation. Specify in advance what would cause you to think the situation may be changing, or that your prior assessment may be wrong. This will alert others to watch for and recognize the potential significance of those developments. It will also help you avoid the common fallacy of rationalizing such developments, if they occur, as not really requiring any modification of your thinking.

Criminal or counterintelligence investigators can use the results to identify the most important leads that require further investigation. By focusing on the most diagnostic or discriminating evidence that appears at the top of the matrix (after using the diagnosticity function), analysts can help focus an investigation on those issues most likely to break the case. This can greatly increase investigative efficiency.

[Contents](#)

Reporting ACH Results

ACH is an analytical tool that generally stays behind the scenes. Your report will usually not mention the ACH analysis unless the report has been tasked specifically as an ACH analysis and is identified as such. Publishing a table listing the Inconsistency Score for each hypothesis or a screen shot showing the matrix would be misleading, suggesting an unwarranted level of precision. You should make the decision on the most likely hypothesis or hypotheses based on your evaluation of the evidence and then craft your narrative justifying your judgment based on the most diagnostic or discriminating evidence.

Your report should, however, reflect the good analytical practices embodied in ACH. The evidence or indicators that have the greatest impact on the Inconsistency Scores are usually cited as key factors in the written analysis. When one recognizes the importance of proceeding by eliminating rather than confirming hypotheses, it becomes apparent that a written argument for a certain judgment is incomplete unless it also identifies alternative judgments that were considered and why they were rejected. The report should communicate the level of uncertainty, the reasons for this uncertainty, any alternative outcomes for which a decisionmaker might need to be prepared, and milestones for future observation that may indicate events are proceeding differently than expected.