

The Consequence Tree Method* for Fact-Supported Decision Making

Jen Schellinck and Tristan Shaeen

Introduction

Is it possible to consistently make better decisions? Is it possible to know if a particular decision was better, or even optimal, relative to other options? These are difficult questions to answer, because once a choice of action has been made, we can't really know what the outcomes of other choices would have been. Humans have an advantage relative to this constraint: we can imagine, both in advance and after the fact, a range of different choices that potentially *could* be taken under the circumstance (or *could have been taken*, once an action is chosen), as well as possible outcomes and consequences of these choices.

When we consider different possible choices and outcomes, we are imagining different possible worlds. In this white paper we present the key ideas and main components of The Consequence Tree Method*, which takes advantage of this ability to imagine possible worlds in order to help people make good decisions in the face of uncertainty.

This method is informed by a range of decision scientist findings, theories and perspectives, including, most significantly:

- Annie Duke (How to Decide: Simple Tools for Making Better Choices)
- Daniel Kahneman (Thinking, Fast and Slow)
- Cassie Kozyrkov (Data Driven? Think Again)

Imagining different possible worlds is central to the method, as is creating a "consequence tree", which captures and maps out the web of possible worlds related to the current situation and available choices. Creating the consequence tree helps to tease apart the good possible worlds from bad ones and map out a path that's more likely to end up in one of the good worlds.

Pulling in or generating relevant facts and creating a "box of facts", or "knowledge base" is key to the successful navigation of possible paths, since this information helps us to better imagine, distinguish and choose between possible worlds.

*(colloquially referred to as the "consequence octopus")

It can also suggest which possible worlds are more or less likely to occur. This is what we mean when we say that the consequence tree method enables fact-supported decision making.

The more relevant evidence there is, and the more that it is used comprehensively and appropriately, the more predictive, accurate and useful the consequence tree and subsequent decision will be.

Using this method requires some time and effort. Whether or not that time and effort is worth it will depend on the importance of the decision you are facing, and the availability of the required resources.

The Structure of the Consequence Tree

For a given situation, consider the possible worlds that could exist because of different choices made, along with the potential outcomes or consequences of each choice. In the context of the Consequence Tree Method, all of these possible choices and consequences grow out from the root of a consequence tree, where the root represents the current situation. Technically speaking, the current situation is the actual world, not a possible world. However, because we have imperfect knowledge about the world we are in, our current situation as we understand it is effectively a possible world too.

Each of the tree's first set of branches represents a particular choice of action, incorporated into a possible “choice world” that contains this action. For example, faced with a splitting path on a hike, one choice world represents walking down the left path, another walking down the right path, and a third, turning around and walking back to the trail head.

Each of these possible choice worlds then becomes the starting point for a new set of branches, representing potential alternative outcomes (along with any side effects) of that particular chosen action. In this second stage, the assumption is that we’re lacking perfect knowledge of the choice world, so different outcomes are possible. In the path example, if we have no map and haven’t hiked the trail before, taking the left path could possibly lead back to the parking lot or it could come to an end at a pretty waterfall. These possible “outcome worlds” may be desirable or undesirable, both with respect to how well they meet a given goal, and with respect to the broader acceptability of their outcomes, impacts and side effects. Each outcome world can be rated relative to their desirability, along with how likely they are to come to pass.

These considerations can be captured in a consequence tree that has some similarities to a decision tree with many branches. However, a consequence tree is not the same as a decision tree. In a decision tree, the different branches represent different properties of a situation or object, and following the branches ultimately leads to tree leaves that

represent a classification of the object or situation. A consequence tree does incorporate and implicitly represent different factors relevant to the decision-making process, but the focus when creating the consequence tree is on imagining and making explicit different possible worlds, rather than detailing the factors and factor values involved in categorizing the different outcomes.

The Consequence Tree Method: Elements and Stages

With some of the high-level details now explained, this section will expand upon the elements and stages involved in creating a consequence tree and then taking a decision. The description here still allows for a variety of different ways to proceed, and although the method is presented in a particular sequence of stages here, these stages can also be carried out in different orders and iterated over multiple times.

Elements considered and incorporated into the consequence tree include:

- Situation + Problem + Goal
- Values and Preferences
- A Default Decision
- Information, Facts, Knowledge
- Decision Factors
- Choices, Possible Actions
- Direct Outputs and Results
- Possible Long-Term Outcomes and Consequences
- Post Decision Evaluation

Defining the Decision Trigger: Goal, Problem or Situation

Kicking off the decision-making process, there will be some situation or event that is responsible for moving you, the decision maker, into an active decision-making state. This might be a goal that you wish to achieve, a problem that must be addressed, a situation that demands action, a question that needs to be answered or even some combination of these. It may be useful to reframe any of the latter (problem, situation, question) in goal terminology before proceeding.

There may be multiple goals at play. In this case, you should define, as much as possible, all of the relevant goals, and then prioritize them. This will determine which ones will have precedence when rating and ranking consequences in the consequence tree. Later steps may prompt a return to, and re-prioritizing of this list, so don't worry too much about making a totally comprehensive list initially.

Describe Values and Considerations

Decisions are supported by values as much as by information. In the Consequence Tree Method, it's important to list and describe your current values and any overarching objectives, policies, mandates and other considerations that must kept in mind and adhered to during and after decision making.

Confirm that the goals as initially defined are consistent with your values and considerations. If not, redefine the goals so that they are consistent with these considerations. In the case of problems, situations or questions, first rephrase them as a goal (e.g., "My goal is to solve the following problem...") and then do this comparison.

When reviewing values and other considerations, take a moment to create value statements that reflect your considerations - i.e. statements like "It's good to...", "It's bad to...", "We want to...", "We want to avoid..."

Add Information on Relevant Topics to the "Fact Box"

While working through the method, you will be generating a fact repository, or "box of facts". You can also think of this as a "knowledge repository", which allows you to get knowledge out of your head and into a place you can review it more explicitly and combine it with other knowledge.

Start the box of facts by adding your value statements. Additional facts will be added over time, such as facts about your current situation and general context, facts that are related to the objects and actions involved in choices you have identified, and facts that are connected to the possible consequences of these choices.

Describe the Current Situation

To create the root of the consequence tree, it's necessary to take stock of the decision-making situation and environment – the decision context. Create a description of the current decision context. Based on this description, add true descriptive statements about the current situation and general decision-making context to your "box of facts". If the current situation is unclear, or difficult to describe, new information might need to be collected to provide a clearer picture of the decision-making context.

Explore Possible Choices

With the current situation described, start to imagine possible choices of action that are consistent with what is currently known about this situation. Try to vividly imagine carrying out each choice, in the process fleshing out the description of the possible “choice world” within which this action is taken. Each of these possible worlds will be added to the consequence tree diagram.

To illustrate this with an example, suppose someone comes to you, holds out a platter with a blueberry muffin and a carrot walnut muffin on it and says, “Would you like a muffin?” You now find yourself in a situation that allows for a choice of action.

As noted above, when deploying the Consequence Tree Method, the first stages for this decision would have been to list any background considerations – values, objectives, policies, etc. – that should be accounted for and then to create a general description of the current situation and decision environment. In the case of the muffin situation, you might have defined ways in which you value your health, or considerations relating to the production of the food. You might have described the situation as one where you are at work in the break room and your co-worker of two years is offering you a muffin.

Once at the stage where you are exploring and defining possible actions – possible choice worlds – you’ll consider the different possible worlds representing choices that can be made in the situation where you’ve been offered the muffin.

In doing this, it will become apparent that these different possible worlds are distinguished from each other by certain aspects that differ from world to world. Depending on the context, people use a variety of terms to capture these differentiators, including:

- Factors
- Considerations
- Aspects
- Options
- Variables
- Influences
- Impacts
- Causes
- Effects

Regardless of the label, it is the specific values of these differentiators which will ultimately determine which choices lead to positive outcomes. They also connect the choices that are available to their particular consequences.

It can be tempting to create a list of choices and choice worlds by first defining a list of differentiating factors, or considerations. For example, in the muffin scenario, differentiating factors or considerations could include the size of each of the muffins, the flavour or ingredients of the muffin and how much you like each muffin. Making a list of such considerations can be helpful when describing the available possible worlds: the world where I pick the big muffin, the world where I pick the small muffin, etc.

At the same time, the number of relevant factors is not the primary focus you need when determining your list of possible worlds. Rather, when defining the possible worlds, the goal is to describe possible worlds that *are exclusive of each other* - i.e., the choices you describe don't overlap with each other, so that if you make one choice, it wouldn't be possible to at the same time make one of the other choices.

Different possible worlds will be described by different collections of statements. Statements that reflect the current, actual world are true statements – i.e. facts. Facts determine which possible world is the real world at any given moment. They also suggest which will become the real world once a choice of action has been made and the action taken.

It's important to keep in mind that possible choices, and by extension choice worlds, are constrained by the current situation (and statements that describe this current situation). If I know that a chocolate chip muffin is not present in the current situation, then the possible world where I choose the chocolate chip muffin is not available to me and doesn't need to be included on in the consequence tree. However, when initially listing choices, facts may be unavailable or information that is thought to be correct may later turn out to be wrong. This is why the Consequence Tree Method lends itself to multiple iterations. As new information comes in, possible worlds that were initially included may be eliminated if the new information indicates they are not, in fact, possible.

You may find that you have groups of possible worlds all falling under a broader description. For example, there might be a group of possible worlds all of which have in common you taking the blueberry muffin off of the platter. Within this broader context, you could then become more and more detailed and specific: there's the possible world where you pick up the blueberry muffin with your left hand, the possible world where you pick up the blueberry muffin with your right hand, etc.

You'll need to exercise your judgement with respect to the level of detail that matters for the decision at hand, as well as which possible worlds would be considered to be relevantly different from each other.

Possible choices of action in the example muffin scenario could include:

- Pick the blueberry muffin
- Pick the carrot walnut muffin
- Pick whichever muffin is the smallest
- Take both muffins off the tray
- Take neither muffin off the tray
- Split the muffins in half and take only part of one or more muffins

As mentioned above, here you must make some decision about the relevant level of granularity appropriate for your list of possible worlds, in order to avoid listing an excessive number of choices. For example, in one possible world you might pick the blueberry muffin and say thank you. In another, you might stand over the platter for five minutes, eventually grab the blueberry muffin off the platter and run out of the room. You will need to decide if any of these details are likely to be relevant, or if "Pick the blueberry muffin" sufficiently captures both of these choices of action.

If desired, you may start with a more complicated description of a particular choice world - for example diagramming out the broader world description and related specific worlds that fit within it - but view this as a draft stage, and conclude this stage by defining a narrowed list of relevant choices that are exclusive of each other.

Note that you can also include choices and possible choice worlds that are not consistent with your values or preferences, or are currently ambiguous on that front. This doesn't mean you might actually choose one of these worlds. Rather, doing this can be beneficial to your overall decision-making process, as it encourages you to fully consider a wide variety of different possibilities. This will be further explored during the consequences stage.

Once you've made your list of possible choices, note which one is your preferred default choice in the absence of any additional information. If you're carrying out this process in a group, it may be worthwhile for people to first note their individual default choice, and then have a group discussion about the default choice. You may also wish to note which choices you are more emotionally drawn to, and which choices your "gut" is telling you to pick.

Explore Consequences of Choices

Once the choice worlds have been determined and described, the next stage is to go through each choice world and consider possible outcome worlds that could unfold from each choice world.

The process for generating possible outcome worlds is similar to that for generating the possible choice worlds. However, instead of considering the choices that are available in the current situation, the approach is shifted to: If this possible choice world were the current situation (the actual world, the possible action described having now been taken, what are the possible results, outcomes, consequences or side effects that could flow from this action.

View this list of possible outcome worlds as a draft that will be further expanded upon and validated at a later stage. As with the choice worlds, depending on the amount of knowledge available now or in the future, it may be that some of the outcome worlds initially thought to be possible will later be ruled as impossible and can be eliminated or adjusted at that point.

Add Still More Information on Consequence-Related Topics to the "Fact Box"

To help with the generation of possible outcome worlds, return to the box of facts. The box should already contain facts which describe values and the current situation. Now add facts to the box that relate to the choices of action that are currently being considered and the possible outcomes of those choices.

In the muffin example, facts in the fact box might include:

- Personal facts: facts about eating preferences, facts about how hungry you are, facts about how well you know the person offering the muffin.
- Food facts: facts about the nutrition composition of the muffins, facts about the possible origin of the ingredients in the muffins
- Health facts: facts about the impact of eating certain types of food on health
- Other facts: facts about how available food is where you are, facts about other people in the room, facts about the person offering the muffin

These facts might originate from previous research, expert knowledge sources or analysis of data.

Sets of facts in the facts box can help clarify and expand upon the initial description of the outcome worlds or even expand the number of outcome worlds. They can be used to add additional details about possible impacts and general outcomes as well as improving an understanding of the immediate consequence of a given action. Alternatively, facts may constrain or alter outcome world descriptions or even entirely eliminate possible worlds if the facts are inconsistent with what has been initially speculated.

Facts that contain statistical information can be used to improve an understanding of how probable (or improbable) each possible outcome world is, relative to the choice of action and other potential consequences being considered. This use of the fact box will be expanded upon in the next stage.

Rate the Probability of Possible Consequences

Throughout the previous steps, you added facts to the facts box about topics pertinent to your situation and choices. Now, review these facts and use them to inform your consequence tree further, as follows:

If statements contradict or further constrain choices of action or possible outcomes of action, you can either eliminate these branches or update them so they are no longer contradictory. For example, suppose in the muffin example, you initially list one of your choices of action as selecting a blueberry muffin, but later learn that you mis-heard the options and there's only a cranberry muffin and a carrot walnut muffin available. In this case update the options and outcomes on the blueberry muffin part of the tree accordingly.

You can also connect pieces of information to other pieces of information. In particular, you can form them into chains of causal statements and then use these chains of statements to both validate and rate the “likely-ness”* of each of the possible outcome worlds. We use the term “likely-ness” here to refer to a more casual notion of how likely an outcome is, avoiding more technical statistical concepts like probability.

A full chain of causal statements looks like the following:

The current situation is Situation A. If Situation A then B, if B then C, if C then Consequence X.

Causal chains can also contain information that is not definitive or deterministic:

The current situation is Situation A. If Situation A then maybe B, if B then maybe C, if C then maybe Consequence X.

Use facts from your fact box to start, fill in or complete causal chains between each possible choice worlds and one or more of the initially generated possible outcome worlds that branch out from that choice world. The creation of casual chains may suggest new outcomes, side effects or other details that had not previously been identified. Consider adding these new outcomes, side effects or details into the consequence tree. Conversely,

if you can't create a casual chain to connect a choice world to an outcome world using information already in the knowledge bank, consider what would have to be the case to create a fully connected causal chain.

Causal chains that deal with non-deterministic (e.g. statistical) facts can be used to assess the likely-ness of a particular outcome world. A simple heuristic to track increases in likely-ness without worrying about quantitative methods could be: starting with a particular choice of action, for every causal chain leading from that action to a particular consequence, add a check mark to that consequence.

As a variant, if one of the facts in the causal chain indicates that the effect and/or consequence is strongly likely, then add more than one checkmark.

With respect to quantitative methods, if you have statistical data relating to the probabilities of certain causal chains, this might also be used to assess likely-ness.

Rate Possible Consequences and Outcomes Relative to Goals, Values, Preferences and Considerations

Look at the possible outcome worlds that have been generated. You've already done some work to evaluate how likely each possible outcome world is, relative to a particular choice of action. Now also rate each of these worlds relative to how well it matches with the goal, problem, situation or question described at the beginning of the exercise. If you are working with a goal, consider how well the outcomes conform to that goal.

If you are working with a problem, consider how well each of the outcomes solves that problem. If you are dealing with a situation that demands action, consider how acceptable the outcomes of action are, relative to your values and other considerations determined. If you are dealing with a question, consider how possible answers align with your considerations.

Eliminate Branches and Take a Decision

Moving backward from the consequences, eliminate any branches of the tree – either consequences or choices - that fall below an acceptable threshold with respect to being inconsistent with values, being too high-risk relative to potential benefits, or which don't adequately match up with the stated goals.

Evaluate the Outcome of the Decision

Once the action has been taken and possible worlds have collapsed into one actual world, with consequences and side effects, close off the decision-making process by reflecting on the outcomes of your choice of action. Did events unfold as expected? Did the anticipated consequences emerge? Is the actual outcome better, worse or just different from what was expected? The “consequence tree” serves as a concrete record of the decision-making process, and allows for a thoughtful engagement with what is vs what could have been. Engaging in this reflection is key for become a better decision maker over time. In this way, the Consequence Tree Method teaches us not just about how the different aspects and outcomes of this particular decision, but also how to become a better decision maker in general.