

The Significance, Persistence, Contingency Framework

William MacAskill, Teruji Thomas (Global Priorities Institute, University of Oxford) and Aron Vallinder (Forethought Foundation for Global Priorities Research)

Global Priorities Institute | September 2022

GPI Technical Report No. T1-2022



The Significance, Persistence, Contingency Framework

William MacAskill*, Teruji Thomas*, and Aron Vallinder†

1 Introduction

The world, considered from beginning to end, combines many different features, or states of affairs, that contribute to its value.¹ The value of each feature can be factored into its *significance*—its average value per unit time—and its *persistence*—how long it lasts. Sometimes, though, we want to ask a further question: how much of the feature’s value can be attributed to a particular agent’s decision at a particular point in time (or to some other originating event)? In other words, to what extent is the feature’s value *contingent* on the agent’s choice? For this, we must also look at the counterfactual: how would things have turned out otherwise?

In this note we give a way to formalise the ideas of significance, persistence, and contingency. We call this the SPC framework. It must be emphasised that the main goal is to help estimate the instrumental value

*Global Priorities Institute, University of Oxford

†Forethought Foundation for Global Priorities Research

This is version 1.1, dated August 2022. We thank Max Daniel for his comments.

¹For these purposes, a ‘feature’ can be any source of value (or disvalue). For example, one could consider a particular headache, the future existence of human civilisation, or the prevalence of QWERTY keyboards. A feature may contribute by having intrinsic value, or by entailing or promoting other things with intrinsic value, such as human welfare.

(typically interpreted as the *expected* value) of an event, compared to the relevant counterfactual. There may be different ways to do this, suitable for different situations. It seems clear to us that thinking in terms of significance, persistence, and contingency can be a useful heuristic when thinking about the importance of historical events and the consequences of our choices for the long-run future; see MacAskill's *What We Owe the Future* for many applications. This note shows how that heuristic can be related to a more formal theory of evaluation. Doing so helps to clarify when the heuristic can be useful; gives some discipline to its application; and opens the door to its use in more formal analyses.

Still, different variations on the SPC framework might be more useful or perspicacious in some situations; in others, the SPC framework may only be useful for evaluating some aspects of an event; and sometimes one may want to use a different framework altogether. So while we hope that this note will be useful to decision-makers, we (of course) make no claim to have the final word.

In section 2, we explain the framework as it appears in *What We Owe the Future*. We'll sketch an alternative approach in section 3. Section 4 discusses how the SPC framework can contribute to the overall evaluation of our options. It explains the 'ITN framework', which evaluates problems in terms of their *importance*, *tractability*, and *neglectedness*, and indicates how the two frameworks can be combined.

2 The Basic Framework

We want to describe the value contributed by some feature F of the world as a result of some originating event E . We will formalise significance, persistence, and contingency in terms of the total value V contributed by F and its duration T , given the occurrence of E . Let ΔV and ΔT be the changes that E makes to V and T , relative to the salient alternative or status quo. Then we can formally define

$$\begin{aligned}
 \text{Sig} &= \Delta V / \Delta T \\
 \text{Per} &= T \\
 \text{Con} &= \Delta T / T.
 \end{aligned}
 \tag{1}$$

The product of these three factors is the change in value:

$$\Delta V = \text{Sig} \times \text{Per} \times \text{Con}. \quad (2)$$

Thus, estimating Sig, Per, and Con leads to an estimate for ΔV . By considering different features of the world, we can then reach an overall evaluation of the event E . (For more on this overall evaluation, see section 4.) Even quite rough estimates may be informative, insofar as we expect some events to differ in value by many orders of magnitude.

While these definitions and the resulting product decomposition of ΔV make logical sense quite generally, they make most intuitive sense in a scenario where the primary effect of E on F is to increase its duration, i.e., to add extra time during which F is present, without changing the value that F contributes during the time it would have been present anyway. Con is high (equal to 1) when the duration of F fully depends on E (F would not have been present but for E), and low (equal to 0) when the event E makes no difference to the duration. More generally, Con tracks the effect of E on the duration of F , and, in the stated scenario, Sig is the average value per unit time contributed during the extra time due to E .

For example, suppose that the value potentially contributed by F at each time increases exponentially, as in the following table:

TIME PERIOD	1	2	3	4
CONTRIBUTION	1	2	4	8

Let us suppose that, in the status quo, F would be present in periods 2 and 3, and the effect of E is to make F present in periods 1 and 4 as well. Then $\text{Per} = 4$, since, given E , F lasts for four periods; $\text{Con} = 1/2$ (half the duration of F is due to E); and $\text{Sig} = 9/2$, the average value contributed per period during the extra two periods attributable to E . The product decomposition (2) gives $\Delta V = 9$. We could, of course, have found this answer directly by adding up the contributions of F in periods 1 and 4.

As this example indicates, Sig differs from both the average value per period in the status quo (i.e., $6/2$) and the average value per period

given E (i.e., $15/4$), though these might both serve as rough approximations. Sig depends on what happens in the status quo as well as what happens given E . It might be conceptually cleaner to have definitions according to which ‘significance’ and ‘persistence’ both depend only on what happens given E , with ‘contingency’ alone taking account of the counterfactual. We will consider such alternative definitions in section 3.

2.1 Marginal Analysis

The originating event E can often be thought of as a marginal change in some parameter X (for example, the level of funding for a particular project). The definitions above then become:

$$\text{Sig} = \frac{dV/dX}{dT/dX} \quad \text{Per} = T \quad \text{Con} = \frac{dT/dX}{T}. \quad (3)$$

The reason we do not simplify Sig to dV/dT is to emphasise that the independent variable is X , which we are assuming has an effect on both V and T as dependent variables.

Note that Con measures the *relative* change in T , that is, the change in T as a proportion of T itself. In the economics literature this would sometimes be called the *semi-elasticity* of persistence. A convenient notation for the semi-elasticity is $\%dT/dX$; thus we can equivalently write

$$\text{Con} = \frac{\%dT}{dX}. \quad (4)$$

At any rate, as the marginal case of equation (2), we have the product decomposition

$$\frac{dV}{dX} = \text{Sig} \times \text{Per} \times \text{Con}. \quad (5)$$

2.2 Uncertainty

The definitions just given can be used in different ways when there is uncertainty about the outcome of E , or about what happens in the

counterfactual. Here, we will assume that the fundamental rule for choice under uncertainty is maximising expected value. So we would like to calculate the change in expected value, $\Delta\mathbb{E}V$, contributed by the possible feature F .

The change in expected value is the same as the expected change in value, $\mathbb{E}\Delta V$. Thus, one possibility is to use the definitions of significance, persistence, and contingency to calculate ΔV in each state, and then to calculate its expectation.

An alternative, which we think will often be more convenient, is to *extend* the definitions of Sig, Per, and Con to take uncertainty into account in each factor:

$$\text{Sig} = \frac{\Delta\mathbb{E}V}{\Delta\mathbb{E}T} \quad \text{Per} = \mathbb{E}T \quad \text{Con} = \frac{\Delta\mathbb{E}T}{\mathbb{E}T}. \quad (6)$$

These ‘ex ante’ interpretations of significance, persistence, and contingency agree with the previous definitions (1) in the absence of uncertainty. We obtain a product decomposition

$$\Delta\mathbb{E}V = \text{Sig} \times \text{Per} \times \text{Con} \quad (7)$$

in the same way as before. The marginal version is exactly similar:

$$\text{Sig} = \frac{d\mathbb{E}V/dX}{d\mathbb{E}T/dX} \quad \text{Per} = \mathbb{E}T \quad \text{Con} = \frac{\%d\mathbb{E}T}{dX} \quad (8)$$

with the result that

$$\frac{d\mathbb{E}V}{dX} = \text{Sig} \times \text{Per} \times \text{Con}. \quad (9)$$

It is worth emphasising that Sig, as defined by (6), is not the same as the *expected* significance, i.e., $\mathbb{E}(\Delta V/\Delta T)$, nor is Con the same as the expected contingency $\mathbb{E}(\Delta T/T)$. Nor is $\Delta\mathbb{E}V$ equal to the product of expected significance, expected persistence, and expected contingency: the expectation of a product is not usually equal to the product of expectations. Similar comments apply to (8) and (9).

Extending our previous discussion, this picture, and especially the definition of Con, makes most intuitive sense when the key margin for

influence is the expected duration, $\mathbb{E}T$. This can involve some combination of increasing the probability that feature F is present and increasing the duration of F , conditional on its presence. In the next section we present an alternative framework that treats these, along with significance, as separate margins for influence.

3 An Alternative Framework

While the framework as described in section 2 captures our intuitions about significance, persistence, and contingency in many situations, we are generally agnostic about what formal framework may be most useful for actual decision-makers, so it may be useful to sketch an alternative which has some conceptual advantages.

First, we will interpret significance simply as the average value contributed by F per unit time, given E :

$$\text{Sig}' = V/T. \quad (10)$$

Persistence will be the same as before:

$$\text{Per}' = T. \quad (11)$$

Thus, note, $V = \text{Sig}' \times \text{Per}'$. Previously, contingency was interpreted as tracking how the duration of F depended on the originating event E ; now, it will track the overall dependence of V on E :

$$\text{Con}' = \frac{\Delta V}{V} \quad \text{or in the marginal case } \frac{\%dV}{dX}. \quad (12)$$

Again, $\%dV$ is short for $\frac{1}{V}dV$. We obtain the usual product decomposition (focusing now on the marginal case):

$$\frac{dV}{dX} = \text{Sig}' \times \text{Per}' \times \text{Con}'. \quad (13)$$

The definitions just given agree with those in section 2 in the special case where F contributes a fixed amount of value per unit time, unaffected by E .

A conceptual advantage of this framework is that significance, like persistence, is understood in a way that is independent of the counterfactual. On the other hand, this makes contingency harder to estimate. However, in the marginal analysis we will have²

$$\text{Con}' = \frac{\%dV}{dX} = \frac{\%d \text{Sig}'}{dX} + \frac{\%d \text{Per}'}{dX}. \quad (14)$$

So we can interpret contingency in terms of the semi-elasticities of significance and of persistence. There may be situations where these semi-elasticities are relatively easy to understand.

What about uncertainty? There are again several ways to proceed. We could use Sig' , Per' , and Con' to calculate dV/dX , and then calculate expectations. Or we can extend the definition of each factor to include uncertainty, as we did in section 2.2, taking $\text{Sig}' = \mathbb{E}V/\mathbb{E}T$, $\text{Per}' = \mathbb{E}T$, and $\text{Con}' = \%d\mathbb{E}V/dX$, which can again be expanded as $\%d \text{Sig}'/dX + \%d \text{Per}'/dX$.

However, in some cases it may be perspicuous to treat the probability of F 's presence as a separate variable, Prob . For example, small changes in the probability that human civilisation continues beyond this millennium are often thought to be a significant source of value.

To do this, let $\mathbb{E}_F V$ be the expected value contributed by F conditional on the presence of F . (It equals $\mathbb{E}V/\text{Prob}$.) Similarly, let $\mathbb{E}_F T$ be the expected duration of F conditional on the presence of F . We then have

$$\mathbb{E}V = \frac{\mathbb{E}_F V}{\mathbb{E}_F T} \times \mathbb{E}_F T \times \text{Prob}. \quad (15)$$

Now define

$$\begin{aligned} \text{Sig}' &= \frac{\mathbb{E}_F V}{\mathbb{E}_F T} = \frac{\mathbb{E}V}{\mathbb{E}T} \\ \text{Per}' &= \mathbb{E}_F T \\ \text{Con}' &= \frac{\%d\mathbb{E}V}{dX} = \frac{\%d \text{Sig}'}{dX} + \frac{\%d \text{Per}'}{dX} + \frac{\%d \text{Prob}}{dX}. \end{aligned} \quad (16)$$

²Here we use the fact that $\%dV/dX = d \log V/dX$; the decomposition $V = \text{Sig}' \times \text{Per}'$ gives $\log V = \log \text{Sig}' + \log \text{Per}'$.

We obtain a four-factor decomposition³

$$\frac{dEV}{dX} = \text{Sig}' \times \text{Per}' \times \text{Prob} \times \text{Con}' . \quad (17)$$

In short, in this alternative framework we treat significance, persistence, and probability as different factors making up the value contributed by F , and define contingency to be the sum of their semi-elasticities.

4 The ITN Framework and Prioritisation

We can embed the SPC framework into the Importance, Tractability, Neglectedness (ITN) framework for prioritising among global problems, which was first proposed by Holden Karnofsky.⁴

The ITN framework is a framework for assessing problems with respect to the marginal benefit of investing work (or other resources) in trying to solve them. The assessment proceeds in terms of three factors. The only one immediately relevant to our discussion is the first:

Importance: how much value there is in solving the problem, or making progress on it. Another word used for the same idea is *scale*.

To formalise this idea, let us make use of a parameter X that measures progress on solving the problem. This could directly reflect the percentage of the problem that has been solved, or it could be some other variable that is a proxy for progress. The importance of marginal progress on the problem (as opposed to the importance of solving it completely) is then naturally identified with

$$\text{Imp} = \frac{dEV_{\text{tot}}}{dX} \quad (18)$$

³An earlier version of this working paper omitted the Prob factor from this equation; thanks to Amane Watahiki for pointing out this mistake.

⁴See Karnofsky (2014), ‘Narrowing down U.S. policy areas’, [link]. A variety of different terminology for the three factors can be found, some of which we’ll mention below.

where V_{tot} is the overall value of the world. Here we have gone directly for a definition of importance that takes into account uncertainty about the value of making progress on the problem, e.g., because of uncertainty about downstream effects. Notice that we do, in principle, want to take into account all downstream effects, hence the use of V_{tot} .

How is this related to the SPC framework? We earlier considered the value V contributed by some feature F of the world, and the change in V attributable to some originating event, such as an increase in the variable X . Since we may want to consider different features of the world, let us now call that value V_F . The role of the SPC framework, in this context, is to help estimate $d\mathbb{E}V_F/dX$ (i.e., Sig \times Per \times Con) for different features F , contributing to an overall estimate of Imp. In the simplest case, we can identify a list of features F_1, F_2, \dots such that

$$V_{\text{tot}} = V_{F_1} + V_{F_2} + \dots . \quad (19)$$

Then Imp is also a sum,

$$\text{Imp} = \frac{d\mathbb{E}V_{F_1}}{dX} + \frac{d\mathbb{E}V_{F_2}}{dX} + \dots . \quad (20)$$

Of course, it may often be difficult to identify a list of easily evaluable features that contribute additively to the value of the world, as in equation (19), although, on the other hand, even rough approximations can be informative.

4.1 Tractability, Neglectedness, and Leverage

We conclude with some further comments about the ITN framework (which are, however, not necessary for understanding the SPC framework).

The second factor used in the ITN framework is

Tractability: how easy it is to solve the problem, or to make progress. Another word used for the same idea is *solvability*.

At a first pass, the more important a problem, and the more tractable it is, the better to work on it; the ideal would be to accomplish something very good with very little effort (one could then spend more effort on other things). Let W be the amount of work done on the problem (or the amount of other resources invested in it). The marginal value of working on the problem is then

$$\frac{dEV_{\text{tot}}}{dW} = \frac{dEV_{\text{tot}}}{dX} \times \frac{dX}{dW}. \quad (21)$$

The first factor on the right-hand side is Imp. It would be natural to identify ‘tractability’ with the second factor, dX/dW , the marginal progress made by extra work. However, it is traditional (and sometimes useful) to separate out of the second factor another consideration:

Neglectedness: how much work is already being done. A word used for the opposite idea is *crowdedness*.

The idea here is that many problems show diminishing returns; that is, if a lot of work has been done, then the easiest aspects of the problem will already have been solved, making it harder to make progress. Owen Cotton-Barratt proposed identifying ‘neglectedness’ with $1/W$, leading to the following decomposition:

$$\frac{dX}{dW} = \frac{dX}{\%dW} \times \frac{1}{W}. \quad (22)$$

He then identifies ‘tractability’ with the first factor on the right-hand side, $dX/\%dW$.⁵

⁵See Cotton-Barratt (2016), ‘Prospecting for Gold’, [link], and also Wiblin (2019), ‘A framework for comparing global problems in terms of expected impact’, [link]. Note that our variable X corresponds formally to Cotton-Barratt’s $\log S$. The issue here is whether we are interested in absolute or relative increases in our measure of progress. For example, if X is the ‘percentage of the problem solved’ then it seems most natural to think in terms of absolute changes in X , as we do in this report. However, it may sometimes be more convenient to focus on some concrete statistic such as (say) ‘number of bednets distributed’. Letting S be this number, it may then be convenient to think in terms of relative changes in S , and thus to define importance as $dEV_{\text{tot}}/\%dS$ (equivalently, $dEV_{\text{tot}}/d \log S$) and tractability as the elasticity $\%dS/\%dW$ (equivalently $d \log S/\%dW$). Thanks to Cotton-Barratt for discussion.

These are perfectly sensible definitions which are useful in many contexts. However, this definition of tractability is arguably most intuitive when there are diminishing returns, and most of all when there are logarithmic returns. This is connected to the point that a high value of $1/W$ only intuitively speaks in favour of additional work if there are diminishing returns. In the simplest case of logarithmic returns, where $X = a \log W$, we find $dX/dW = a/W$. Although, in one sense, we might want to say that the problem becomes ‘less tractable’ as W increases (there are diminishing returns), the constant scale factor a is related to the *overall* difficulty of the problem—the total work that would be required to solve it. In contrast, consider a case in which there are linear returns: $X = W$. Then $dX/dW = 1$. It could be somewhat counter-intuitive to think of tractability increasing with W in this case, and the current value of W does not give a sense of the overall tractability of the problem, either.

An alternative approach is to identify ‘tractability’ with the overall difficulty of the problem. Let W_0 be the total amount of work that would be required to solve the problem (or to achieve some relevant benchmark), and let X_0 be the value of X that would count as a full solution. (The simplest way to think about this is that X is ‘the percent solved’, so that $X_0 = 100$). We can then write

$$\frac{dX}{dW} = \frac{X_0}{W_0} \times \frac{dX/dW}{X_0/W_0}. \quad (23)$$

The first factor on the right-hand side, X_0/W_0 , measures the overall easiness of solving the problem, while the second, $(dX/dW)/(X_0/W_0)$, measures how much easier it is to make progress at the current margin. In other words, it captures directly the extent to which there are lower-than-average hanging fruit. To avoid confusion, let us call this second factor *leverage*. Leverage is a form of ‘neglectedness’ (the low-hanging fruit haven’t been picked), but high leverage does not imply that little work is being done. When there are diminishing returns, leverage decreases with W (for fixed values of X_0 and W_0), but when (for example) returns are linear, leverage is constant: the marginal worker has as much impact as the average one.