



Responsibility Under Uncertainty: Which Climate Decisions Matter Most?

Nicola Botta^{1,2} · Nuria Brede^{1,7} · Michel Crucifix³ · Cezar Ionescu⁴ · Patrik Jansson² · Zheng Li^{5,6} · Marina Martínez³ · Tim Richter⁷

Received: 23 November 2021 / Accepted: 23 October 2022 / Published online: 2 February 2023
© The Author(s) 2023

Abstract

We propose a new method for estimating how much decisions under monadic uncertainty matter. The method is generic and suitable for measuring responsibility in finite horizon sequential decision processes. It fulfills “fairness” requirements and three natural conditions for responsibility measures: agency, avoidance and causal relevance. We apply the method to study how much decisions matter in a stylized greenhouse gas emissions process in which a decision maker repeatedly faces two options: start a “green” transition to a decarbonized society or further delay such a transition. We account for the fact that climate decisions are rarely implemented with certainty and that their consequences on the climate and on the global economy are uncertain. We discover that a “moral” approach towards decision making — doing the right thing even though the probability of success becomes increasingly small — is rational over a wide range of uncertainties.

Keywords Climate policy · Responsibility measures · Uncertainty · GHG emissions processes · Verified decision making

1 Introduction

When a person performs or fails to perform a morally significant action, we sometimes think that a particular kind of response is warranted. Praise and blame are perhaps the most obvious forms this reaction might take. For example, one who encounters a car accident may be regarded as worthy of praise for having saved a child from inside the burning car, or alternatively, one may be regarded as worthy of blame for not having used one’s mobile phone to call for help. To regard

such agents as worthy of one of these reactions is to regard them as responsible for what they have done or left undone [1].

The quote from “The Stanford Encyclopedia of Philosophy” (SEP) provides a compelling account of what responsibility is about. The car accident example is pointed because of two reasons.

First, because it rests on an implicit and widely accepted understanding of what a person “who encounters a car accident” with a child “inside the burning car” shall do. Namely, their best to rescue the kid.

Second, because what the person is regarded as worthy of praise or of blame for having done (or left undone) are *best* and *worst* actions with respect to the goal of rescuing the child: the agent can expect little praise for having used the mobile phone to call for help and possibly also little blame for not having managed to get the child out of the burning car. By contrast, they can expect blame for not having used the mobile phone to call for help.

1.1 Responsibility in Climate Decisions

In the context of climate policy, the measure by which praise and blame shall be attributed to decisions is not always as clear as in the [1] example. This is because of two reasons.

✉ Nicola Botta
botta@pik-potsdam.de

¹ Potsdam Institute for Climate Impact Research, Potsdam, Germany

² Chalmers University of Technology, Göteborg, Sweden

³ Université Catholique de Louvain, Ottignies-Louvain-la-Neuve, Belgium

⁴ Technische Hochschule Deggendorf, Deggendorf, Germany

⁵ Northeastern University, Toronto, Canada

⁶ Arima Inc., Toronto, Canada

⁷ Potsdam University, Potsdam, Germany

We know that decisions that are taken (or delayed) now and in the next decades, e.g., on greenhouse gas (GHG) emissions, are crucial for events that unfold in the centuries and millennia to come, mainly because the physical and chemical processes involved in reabsorbing atmospheric CO₂ are very slow [2].

We also know that current climate policies may lead future generations to (attempt to) mitigate the negative effects of climate change (CC) by adopting geo-engineering measures (for example, massive injections of aerosols in the atmosphere [3–5] that can have severe collateral effects (for instance, on agricultural yield, hydrological events, public health [6], ecosystems [7] or precipitation [8, 9]) or otherwise face enormous human and economic costs.

But, in contrast to decision problems in technical sciences and in engineering, in which the *goal* of decision making is typically well understood, there is little agreement on how to value (and to discount) the chances and the risks of climate change [10].

This is especially true when such risks and the associated potential costs are related to events that unfold in hundreds or thousands of years and thus very much depend on assumptions about the preferences of future generations.

Because of these difficulties, most attempts at estimating the impact of current and near term climate policies are based on comparisons of costs and benefits over a time horizon of one or two hundred years [4, 11–13] and, even so, are controversial [10, 14].

In short: in climate policy one cannot rely on a widely accepted understanding of what the goals of decision making are. Thus, specifying such goals is not as straightforward as in the car accident example.

The second reason why attributing praise and blame to climate decisions is not as straightforward as in the car accident example [1] is uncertainty. Can an agent be held responsible for (performing or for failing to perform) actions that matter very little? What does it mean precisely for decisions to matter?

In the scientific community but also in part of the civil society (think of the “Fridays for future” movement), there is a strong concern that decisions that are taken (or delayed) now will have severe consequences on the options that will (not) be available to upcoming generations.

But what do we mean when we say that current decisions matter more than decisions that will be taken by future generations? Are there systematic ways to measure how much decisions matter when these have to be taken under *epistemic* but also political and social *aleatoric* [10] uncertainty? Is there a natural way of comparing similar decisions at different points in time? To provide accountable confidence that all the efforts that are associated with the actual implementation of such decisions (often involving politically difficult negotiations, changes in legislations, taxation

and incentivations schemes, not to mention technological research and development) are (not) devoted to decisions that (do not) really matter?

1.2 What this Paper is About

We propose a method for measuring how much decisions under uncertainty matter and apply it to a stylized GHG emissions decision process.

The method is an application of the computational theory of policy advice and avoidability originally proposed in [15]. This theory supports the specification of time-discrete sequential decision problems [16, 17] and the computation of verified *best* decisions under uncertainty. It is an extension of the formal framework of vulnerability [18] and of the notion of *monadic* dynamical system originally introduced in [19] and allows dealing with different kinds of uncertainty in a logically consistent manner. The theory is formulated in Idris [20, 21], an implementation of type theory [22].

For the sake of providing a self-contained account of our method, we summarize the elements of the theory [15] that we apply in this work in the next section. Readers familiar with [15] can skip Section 2 and jump directly to Section 3. The [15] theory is formulated in Idris, a dependently typed functional language [20, 21]. Many climate scientists are well acquainted with imperative languages but less so with functional, dependently typed languages. For these readers, we provide a minimal introduction to the notation applied in Sections 2 to 6 in the Appendix.

The method for measuring how much decisions under uncertainty matter is based on the observation that many processes in which decisions have to be taken sequentially and under uncertainty can be represented by finite decision networks. We introduce finite decision networks formally in Section 2. Intuitively, a finite decision network is a network in which each decision yields a finite number of possible outcomes.

Because the car accident example from [1], the stylized decision process outlined in Section 1.3, and many interesting decision processes in climate policy can all be represented as finite decision networks [23], we can apply the theory of Section 2 to study such processes.

In particular we can apply the theory to compute a *best* and a *worst* decision at each node (decision step) of the network: The idea is then to measure how much decisions matter by comparing the values (for a specific decision-making goal) associated with such best and worst decisions. If the values of best and worst decisions turn out to be the same, then decisions at that decision step do not matter. By contrast, the larger the difference between the value of best and worst decisions, the more decisions do matter.

In Section 4, we formulate this simple principle and define a measure of how much decisions matter for the

stylized decision process of Section 1.3. The process is formally specified in Section 3.

In Section 5, we extend the specific measure of Section 4 to generic responsibility measures. This is done by introducing a small domain-specific language (DSL) for expressing decision-making goals (in the car accident example, rescuing the child), measures of uncertainty and methods for computing differences in the value of decisions.

These generic responsibility measures account for all the knowledge which is encoded in a sequential decision process or network. They are agnostic with respect to both decision makers and decision steps: how much a decision matters does not depend on the aims or on the preferences of the (real or hypothetical) decision maker; all decisions are measured in the same way. The conditions ensure that responsibility measurements are *fair*. Here, “fair” is a technical notion and we make no claims about fairness in any wider sense. In particular, this technical notion shall not be confused with that of generational fairness discussed in Section 7.

Before outlining the GHG emissions process that we will use to illustrate our method for measuring how much decisions under uncertainty matter, let us discuss a potential criticism.

We have explained that we measure how much decisions matter at a given decision step (under uncertainty about the consequences of such decisions both at that step and at future steps) by applying the theory of Section 2 to compute best and worst decisions. What is the added value of measuring how much decisions matter for policy making if we already know how to take best decisions? This is a legitimate and important question to which we want to provide a first answer right now.

Remember that, in order to obtain best (and worst) decisions for a specific decision step one has to specify a goal of decision making. For example, rescuing the child in the car accident example or, as we will see in the next section, avoid long term climate change impacts or short term economic downturns.

The value of best and worst decisions and thus how much decisions matter will then typically be different for different goals. Best decisions under a given goal might be suboptimal (or even worst) under another goal. Decisions that matter a lot for a given goal might turn out to be irrelevant for another one.

In Section 1.1 we have pointed out how difficult and controversial it is to specify such goals in climate decision processes, see also the discussion on the impossibility of “value-free” climate science in [10]. Thus, the added value of our measures is that of providing a better understanding of how the importance of specific decisions depends on the (possibly conflicting) goals of decision-making and also on the measure of uncertainty (expected value, worst-case

value) and on other aspects of sequential decision processes discussed in Section 2

Our hope is that understanding that a specific climate policy (say, pushing forward a “green” transition right now) may be crucial or irrelevant depending on which measures of uncertainty and goals are put forward for the decision process at stake will lead to a more rational and collaborative approach, for example in climate negotiations.

Thus, besides proposing a novel approach to the problem of rational choice and attribution of responsibility [24–26], our work is a contribution to pragmatic decision making under uncertainty with a specific focus on climate decisions.

1.3 A Stylized Decision Process

Consider a GHG emissions process in which now and for a few more decades, humanity (taken here as a global decision maker) faces two options:

1. Start a “green” transition by reducing GHG emissions according to a “safe” corridor, for example, the one depicted at page 15, Figure SPM.3a of the IPCC Summary for Policymakers [27]
2. Delay such transition.

In other words, assume that, over the time period between two subsequent decisions (say, for concreteness, a decade), either a transition to a nearly decarbonized global socio-economic system is started or nothing happens. Further, assume that, once a transition has been started, it cannot be halted or reversed by later decisions or events. We consider this oversimplified situation only for the sake of clarity, although it might well be that green transitions are in fact fast and irreversible [28].

Selecting to start a green transition in a specific physical, social and economic condition yields a different “new” condition at the next decision step. Let’s call one such condition a *micro-state*.

The idea is that micro-states are detailed descriptions of physical, social and economic observables. For example, a micro-state could encode values of GHG concentrations in the atmosphere, carbon mass in the ocean upper layer, global temperature deviations, frequency of extreme events, values of economic growth indicators, measures of inequality, etc. Even if we knew the “current” micro-state perfectly, the set of possible micro-states at the next decision step (say, one decade later) would still be extremely large, reflecting both the epistemic uncertainties (imperfect knowledge) about the (physical, social and economical) processes that unfold in the time between now and the next decision step and the aleatoric uncertainty [10] of those processes.

Descriptions of decision processes explicitly based on micro-states would be both computationally intractable

and, as discussed in detail in Section 3.3, methodologically questionable. As in the car accident example quoted at the opening of this section, we avoid these shortcomings by considering only a small number of sets (clusters, partitions) of micro-states. These *macro-states* (in the following, just states) consist of micro-states in which:

- A green transition has been *started* or *delayed* (*S*-states, *D*-states).
- The economic wealth is *high* or *low* (*H*-states, *L*-states).
- The world is *committed* or *uncommitted* to severe CC impacts (*C*-states, *U*-states).

In other words, we only distinguish between 8 possible states: *DHU*, *DHC*, *DLU*, *DLC*, *SHU*, *SHC*, *SLU* and *SLC* where *DHU* represent micro-states in which a green transition has been *delayed*, economic wealth is *high* and the world is *uncommitted* to future severe CC impacts. Similarly for *DHC*, *DLU*, etc.

Clearly, this is a very crude simplification. But it is useful to study the impact of uncertainty on relevant climate decisions and sufficient to illustrate our approach towards measuring how much decisions matter. Also, notice that binary partitioning of micro-states is at the core of the original notion of planetary boundaries [29], of the topological classification proposed in [30] and of the social dilemmas discussed in [31].

The decision process starts in *DHU*. In this state, a decision to start a green transition can lead to any of the *DHU* ... *SLC* states, albeit *with different probabilities*: the idea is that the probability of reaching states in which the green transition has been started (*S*-states) is *higher* than the probability of reaching *D*-states, in which the green transition has been delayed. Symmetrically, we assume that the decision to delay the start of a green transition in *DHU* is more likely to yield *D*-states than *S*-states.

In other words, we assume our (global, idealized) decision maker to be *effective*, but only to a certain degree. This accounts for the fact that, in practice, decisions are not always implemented, be this because global coordination is necessarily imperfect, because global players tend to be in competition and legislations tend to have large inertia or perhaps because some other global challenge (a pandemic or an economic downturn) has taken center stage. As demonstrated in [32], limited effectiveness has a significant impact on optimal GHG emissions policies. Thus, it would be inappropriate to assume that decisions are always implemented with certainty.

Another essential trait of our stylized process is that decisions to start a green transition, if implemented, are more likely to yield states with a low level of economic wealth (*L*-states) than states with high economic wealth. This assumption reflects the fact that starting a green transition requires

more investments and costs than just moving to states in which most of the work towards a globally decarbonized society remains to be done.

Finally, we assume that the probability of entering states in which the world is committed to severe CC impacts is higher in states in which a green transition has not already been started as compared to states in which a green transition has been started. Also, as one would expect, delaying transitions to decarbonized economies *increases* the likelihood of entering states in which the world is committed to severe CC impacts.

We give a complete formal specification of our stylized decision process in Section 3. Before turning to Section 2, let's look a bit more closely at the notion of responsibility discussed so far.

1.4 Clarifications, Caveats and Related Work

The notion of responsibility illustrated by the car accident example depends on a number of factors.

First and foremost, we have an entity capable of taking decision: the “one who encounters a car accident”. In the stylized decision process outlined in Section 1.3, we have referred to this entity as to the *decision maker*.

Second, we have situations like “encountering a car accident” or like “the child being saved”. These are coarse, macroscopic descriptions of initial, intermediate or final stages of a decision process that unfolds in time. In our stylized decision process, we have used the term *state* to denote such coarse descriptions or, more concretely, sets of possible micro-states. We formalize the notion of state in Section 2.

The third important element we have is *options*. In the SEP example, the decision maker (the “one who encounters a car accident”) may “be regarded as worthy of praise” or “may be regarded as worthy of blame” for having or for not having used a mobile phone to call for help: decision makers have to be capable of performing certain actions (using a mobile phone to call for help, save the child) for being “regarded as responsible for what they have done or left undone”.

In our stylized decision process we have maintained that, in the initial state *DHU*, the decision maker is, up to a certain extent, capable of starting a green transition or to delay it. In this state, they might be held responsible for having or for not having started the transition.

Notice that the options available to the decision maker in a given state typically depend on that state and, in general, also on the point in time (decision step) at which that state has been obtained.

Also notice that, while the *for* in “for having or for not having” is relative to a decision taken, the praise or the blame and therefore the extent to which the decision maker is regarded as responsible crucially depend on the

consequences of that decision with respect to a goal: the child being saved, the economic wealth being high, the world not being committed to severe CC impacts.

A few remarks are in order here. First, notice that such a goal does not need to relate to what the decision maker considers to be desirable or worth pursuing in the decision process at stake.

Second, for a decision maker to be held responsible for a decision in a given state, say x , the goal of decision making has to be specified in terms of future states that are obtainable from x . If this is not the case, the decisions to be taken in x are not *causally relevant* [33] and any responsibility measure should return a verdict of “not responsible”.

Finally, a necessary condition for a decision maker to be held responsible for a decision in a given state is that at least two choices are available in that state, under the principle that one cannot be held responsible if one has no choice.

We conclude this introduction with a few caveats. The first one is about the notion of responsibility itself: there is a huge literature on the problem of measuring (quantifying, attributing, etc.) responsibility.

Common approaches distinguish at least between *ex-ante* and *ex-post* notions of responsibility [34] and, when more entities contribute to a decision, for example in voting schemes or international agreements, between *individual* and *collective* responsibility [33].

In the context of law, notions of ex-post responsibility are crucial, e.g., to quantify liability for harm. But for the kind of GHG emissions decision processes exemplified by our stylized process and as a guideline for decision making, ex-ante responsibility is the relevant notion.

Another caveat is about the notion of stylized decision process itself. We have introduced this notion in [32] and we will discuss it in more detail in Section 3.3. The notion is closely related with that of *storyline* put forward in [10] but there are also important differences. The storyline approach has been proposed to overcome the (essentially unavoidable) ineffectiveness of predictions of climate change impacts at regional scales. It maintains that, at such scales, questions of climate risks (for given scenarios) need to be reframed “from the ostensibly objective prediction space into the explicitly subjective decision space”. The distinction between epistemic and aleatoric uncertainty and the “identification of physically self-consistent, plausible pathways” are pivotal for such reframing and “the mathematical framework of a causal network” is the key for “reconciling storyline and probabilistic approaches”.

The notion of stylized decision process accounts for the fact that at the global scale “climate decisions are not made on the basis of climate change alone”, are rarely implemented with certainty and can easily be sidetracked by other global challenges, as discussed in Section 1.3. As a consequence, questions of climate policy need to be studied in

“the explicitly subjective decision space” at both the global and the local scale.

As in the storyline approach put forward in [10], the key for applying stylized decision processes is a mathematical framework. In our case, this is provided by the theory [15], and the causal networks proposed in [10] are a special case of decision networks, see also Sections 2 and 3. To the best of our knowledge, [15] is still the only theory for computing optimal policies for decision making under monadic uncertainty that has been *verified*. This means that the policies obtained with the theory can be machine-checked to be optimal. The possibility of computing verified optimal policies was one of the two main motivations (the other one being the capability of enforcing transparency of assumptions) for formulating the theory in a dependently typed language. A consequence of this is that the best and the worst decisions that define the responsibility measures proposed in our application are provably best and the worst decisions. We believe that providing this level of guarantees is crucial in climate decision making: in contrast to policy advice in, e.g., engineering and logistics, recommendations to decision makers in matters of climate policy cannot undergo empirical verification. Thus, in climate policy advice, the only guarantees that advisors can provide to decision makers have to come from formal methods and verified computations, which is the highest standard of correctness that science can provide today.

A final caveat is about what this paper is not about. We develop a formal method to understand which decisions under uncertainty matter most and apply this method to a decision problem of *global* climate policy. Our aim is neither to recommend climate actions nor to design specific mechanisms, e.g., to improve coordination and collaboration between decision makers. First and foremost, we aim at better understanding climate decision making under uncertainty.

2 The Theory in a Nutshell

In this section, we overview the elements of the [15] theory¹ that we apply in Sections 3 to 6. For motivations, comparisons with alternative formulations and details, please see [15, 35]. For a summary of the notation, see the [Appendix](#).

In a nutshell, the theory consists of two sets of components: one for the *specification* of sequential decision problems (SDPs) and one for their *solution* with verified backward induction. For informal introductions to SDPs, see [15]. Reference mathematical introductions to SDP are given in sections 1.2 and 2.1 of [17] and [16], respectively.

¹ Hereafter simply referred to as the *theory*

For an application of the theory to GHG emissions problems, see [32].

The components for the specification of SDPs are global declarations. Four of these describe the sequential decision process that underlies a decision problem. The first declaration

$$M : \text{Type} \rightarrow \text{Type}$$

specifies the uncertainly monad M . Discussing the notion of monad here would go well beyond the scope of this manuscript, and we refer interested readers to [36] and [37]. The idea is that M accounts for the uncertainties that affect the decision process. In the stylized GHG emissions process outlined in the introduction, M represents stochastic uncertainty. For this process, values of type MA are finite probability distributions on A , see Section 3.

Remember that, as shown in Section 1.3, sequential decision processes are defined in terms of the states, of the options available to the decision maker (in a given state and at a given decision step) and of the state transitions that take place between two subsequent decisions.

In control theory, the options available to the decision maker are called *controls* and the theory supports the specification of the *states*, of the *controls* and the state *transition function* of a decision process in terms of three declarations:

$$X : (t : \mathbb{N}) \rightarrow \text{Type}$$

$$Y : (t : \mathbb{N}) \rightarrow X t \rightarrow \text{Type}$$

$$\text{next} : (t : \mathbb{N}) \rightarrow (x : X t) \rightarrow Y t x \rightarrow M (X (S t))$$

The interpretation is that $X t$ is the type (set) of states at decision step t . For example, the states DHU, DHC, \dots, SLC of our stylized GHG emission process. Similarly, $Y t x$ represents the controls available at decision step t and in state x and $\text{next } t x y$ is an M -structure of the states that can be obtained by selecting control y in state x at decision step t . In the decision process of Section 1.3, $Y 0 DHU$ (the set of controls available to the decision maker at decision step 0 and in state DHU) only contains two alternatives: *Start* and *Delay*.

The uncertainty monad M , the states X , the controls Y and the transition function next completely specify a decision process: if we were given a rule for selecting controls for a given decision process (that is, a function that gives us a control for every possible state) and an initial state (or, in case of epistemic uncertainty [10], a probability distribution of initial states) we could compute all possible trajectories compatible with that initial state (or with that probability distribution) together with their probabilities².

Indeed, a sequential decision *problem* for n steps consists of finding a sequence of n *policies* (in control theory, functions that map states to controls are called policies) that, for a given decision process, maximizes the *value* of taking n decision steps according to those policies, one after the other.

Here, the value of taking n decision steps according to a sequence of n policies is defined through a measure (in stochastic problems often the expected-value measure) of a sum of rewards obtained along the trajectories. It follows that, in order to fully specify a decision problem, one has to define the rewards obtained at each decision step, the sum that the decision maker seeks to maximize and the measure function. In the [15] theory, this is done in terms of 6 problem specification components. These are summarized in the next section.

2.1 Problem Specification Components

$$\text{Val} : \text{Type}$$

$$\text{reward} : (t : \mathbb{N}) \rightarrow (x : X t) \rightarrow Y t x \rightarrow X (S t) \rightarrow \text{Val}$$

$$(\oplus) : \text{Val} \rightarrow \text{Val} \rightarrow \text{Val}$$

Here, Val is the type of rewards, $\text{reward } t x y x'$ is the reward obtained by selecting control y in state x when the next state is x' and the infix operator \oplus is the rule for adding rewards. A few remarks are at place here.

1. In many applications, Val is a numerical type and controls are actions that consume certain amounts of resources: fuel, water, etc. In these cases, the reward function encodes the value (cost) of these resources (and perhaps also the benefits achieved by using them) over a decision step. Often, the latter also depends on the “current” state x and on the next state x' . For example, in the stylized decision problem of Section 1.3, $\text{reward } t x y x'$ would possibly be higher than $\text{reward } t x y x''$ if x' is an H-state (a state with a high level of economic wealth) and x'' is an L-state. The theory nicely copes with all these situations.
2. When Val is a numerical type, \oplus is often the canonical addition associated with that type. However, in many applications more flexibility is needed, e.g., to account for the fact that later rewards are often valued less than earlier ones. Again, formulating the theory in terms of a generic addition rule nicely covers all these applications.
3. Mapping $\text{reward } t x y$ onto $\text{next } t x y$ ³ yields a value of type $M \text{Val}$. These are the *possible* rewards obtained

² This is not a trivial result. It holds because we have required M to be a monad.

³ Because M is a monad, functions of type $A \rightarrow B$ can be mapped on values of type MA , obtaining values of type MB for arbitrary $A, B : \text{Type}$

by selecting control y in state x at decision step t . A sequential decision problem for n steps consists of finding a sequence of n policies that maximizes a measure of a sum of the rewards along possible *trajectories*. We introduce a *value function* that computes such a measure in Section 2.2: as it turns out, comparing two policy sequences for a fixed initial state essentially means comparing two M *Val* values.

In mathematical theories of optimal control, the implicit assumptions are often that *Val* is equal to \mathbb{R} , values of type M *Val* are probability distributions on real numbers and such values are compared in terms of their *expected value* measures. Measuring uncertainties in terms of expected value measures subsumes a neutral attitude towards risks. This is not always adequate and the theory supports alternative (e.g., worst-case) measures via the declaration:

$$meas : M \text{ Val} \rightarrow \text{Val}$$

In much the same way, the framework allows users to compare *Val* values in terms of a problem-specific total preorder

$$\begin{aligned} (\leq) & : \text{Val} \rightarrow \text{Val} \rightarrow \text{Type} \\ \text{lteTP} & : \text{TotalPreorder} (\leq) \end{aligned}$$

This allows, among others, to specify multi-objective optimal control [13] problems. Here \leq and *Total Preorder* : $(A \rightarrow A \rightarrow \text{Type}) \rightarrow \text{Type}$ are predicates like those discussed in Section A8 and *TotalPreorder* R encodes the notion that R is a total preorder.

2.2 Problem Solution Components

The second set of theory components formalizes classical optimal control theory. Here, we only provide a concise, simplified overview. Motivation and details can be found in [38, 15] and [32]. For an introduction to the mathematical theory of optimal control, we recommend [16] and [17]. As mentioned, policies (decision rules) are functions from states to controls:

$$\begin{aligned} \text{Policy} & : (t : \mathbb{N}) \rightarrow \text{Type} \\ \text{Policy} t & = (x : X t) \rightarrow Y t x \end{aligned}$$

Policy sequences of length $n : \mathbb{N}$ are then just vectors (remember Section A7) of n policies:

$$\begin{aligned} \text{data PolicySeq} & : (t : \mathbb{N}) \rightarrow (n : \mathbb{N}) \rightarrow \text{Type} \text{ where} \\ \text{Nil} & : \{t : \mathbb{N}\} \rightarrow \text{PolicySeq } t \text{ } Z \\ (\cdot : \cdot) & : \{t, n : \mathbb{N}\} \rightarrow \text{Policy } t \rightarrow \text{PolicySeq } (S t) n \rightarrow \\ & \text{PolicySeq } t (S n) \end{aligned}$$

Perhaps, the most important notion in the mathematical theory of optimal control is that of *value function*. The value

function takes two arguments: a policy sequence ps for making n decision steps starting from decision step t and an initial state $x : X t$. It computes the value of taking n decision steps according to the policies ps when starting in x :

$$\begin{aligned} \text{val} & : \text{Functor } M \Rightarrow \{t, n : \mathbb{N}\} \rightarrow \text{PolicySeq } t n \rightarrow X t \rightarrow \text{Val} \\ \text{val } \{t\} \text{ Nil } x & = \text{zero} \\ \text{val } \{t\} (p :: ps) x & = \text{let } y = p x \text{ in} \\ & \text{let } mx' = \text{next } t x y \text{ in} \\ & \text{meas } (\text{map } (\text{reward } t x y \oplus \text{val } ps) mx') \end{aligned}$$

Notice that, independently of the initial state x , the value of the empty policy sequence is *zero*. This is a problem-specific reference value

$$\text{zero} : \text{Val}$$

that has to be provided as part of the problem's specification. The value of a policy sequence consisting of a first policy p and of a tail policy sequence ps is defined inductively as the measure of an M -structure of *Val* values. These values are obtained by first computing the control y dictated by p in x , the M -structure of possible next states mx' dictated by next and finally by adding $\text{reward } t x y x'$ and $\text{val } ps x'$ for all x' in mx' . The result of this functorial mapping is then measured with the problem-specific measure $meas$ to obtain a result of type *Val*. The function which is mapped on mx' is just a lifted version of \oplus , as one would expect:

$$\begin{aligned} (\oplus) & : \{A : \text{Type}\} \rightarrow (f, g : A \rightarrow \text{Val}) \rightarrow A \rightarrow \text{Val} \\ f \oplus g & = \lambda a \Rightarrow f a \oplus g a \end{aligned}$$

As shown in [35], $\text{val } ps x$ does indeed compute the $meas$ -measure of the \oplus -sum of the *reward*-rewards along the possible trajectories starting at x under ps for sound choices of $meas$. The advantage of the above formulation of val [16, 17, 39] is that it can be exploited to compute policy sequences that are provably optimal in the sense of

$$\begin{aligned} \text{OptPolicySeq} & : \text{Functor } M \Rightarrow \{t, n : \mathbb{N}\} \rightarrow \\ & \text{PolicySeq } t n \rightarrow \text{Type} \\ \text{OptPolicySeq } \{t\} \{n\} ps & = (ps' : \text{PolicySeq } t n) \rightarrow \\ & (x : X t) \rightarrow \\ & \text{val } ps' x \leq \text{val } ps x \end{aligned}$$

Notice the universal quantification in the definition of *OptPolicySeq*: a policy sequence ps is said to be optimal iff $\text{val } ps' x \leq \text{val } ps x$ for any ps' and for any x . The generic, verified implementation of backward induction from [15] is a simple application of Bellman's principle of optimality, often referred to as Bellman's equation [39]. It can be suitably formulated in terms of the notion of *optimal extension*. A policy $p : \text{Policy } t$ is an optimal extension of a policy sequence $ps : \text{Policy } (S t) n$ if it is the case that the value

of $p :: ps$ is at least as good as the value of $p' :: ps$ for any policy p' and for any state $x : X t$:

$$\begin{aligned} \text{BestExt} : \text{Functor } M \Rightarrow \{t, n : \mathbb{N}\} \rightarrow \\ \text{PolicySeq } (S t) n \rightarrow \text{Policy } t \rightarrow \text{Type} \\ \text{BestExt } \{t\} ps p = (p' : \text{Policy } t) \rightarrow \\ (x : X t) \rightarrow \text{val } (p' :: ps) x \leq \text{val } (p :: ps) x \end{aligned}$$

With this formalization of the notion of optimal extension, Bellman's principle can then be formulated as

$$\begin{aligned} \text{Bellman} : \text{Functor } M \Rightarrow \{t, n : \mathbb{N}\} \rightarrow \\ (ps : \text{PolicySeq } (S t) n) \rightarrow \text{OptPolicySeq } ps \rightarrow \\ (p : \text{Policy } t) \rightarrow \text{BestExt } ps p \rightarrow \\ \text{OptPolicySeq } (p :: ps) \end{aligned}$$

In words: *extending an optimal policy sequence with an optimal extension (of that policy sequence) yields an optimal policy sequence*. Another way of expressing the same principle is to say that prefixing with optimal extensions preserves optimality. Proving Bellman's optimality principle is almost straightforward and crucially relies on \leq being reflexive and transitive (remember that \leq is a total preorder). With *Bellman* and provided that we can compute best extensions of arbitrary policy sequences

$$\text{bestExt} : \text{Functor } M \Rightarrow \{t, n : \mathbb{N}\} \rightarrow \text{PolicySeq } (S t) n \rightarrow \text{Policy } t$$

$$\begin{aligned} \text{bestExtSpec} : \text{Functor } M \Rightarrow \{t, n : \mathbb{N}\} \rightarrow \\ (ps : \text{PolicySeq } (S t) n) \rightarrow \text{BestExt } ps (\text{bestExt } ps) \end{aligned}$$

it is easy to derive a verified, generic implementation of backward induction:

$$\begin{aligned} \text{bi} : \text{Functor } M \Rightarrow (t : \mathbb{N}) \rightarrow (n : \mathbb{N}) \rightarrow \text{PolicySeq } t n \\ \text{bi } t Z = \text{Nil} \\ \text{bi } t (S n) = \text{let } ps = \text{bi } (S t) n \text{ in } \text{bestExt } ps :: ps \end{aligned}$$

For this implementation, a machine-checked proof that *bi t n* is an optimal policy sequence for any initial time t and number of decision steps n :

$$\begin{aligned} \text{biLemma} : \text{Functor } M \Rightarrow (t : \mathbb{N}) \rightarrow (n : \mathbb{N}) \rightarrow \\ \text{OptPolicySeq } (\text{bi } t n) \end{aligned}$$

is a straightforward computation, see [15, 35].

2.3 Theory Wrap-up

The components discussed in the last two sections are all what is needed to define the measures of how much decisions matter that we have discussed in the introduction. We introduce these measures in Sections 4 and 5. As discussed in [35], the [15] theory is slightly more general (but also more difficult

to apply) than the one summarized above. The price that we have to pay for the simplifications introduced here are two additional requirements. First, controls have to be non-empty:

$$\text{notEmptyY} : (t : \mathbb{N}) \rightarrow (x : X t) \rightarrow Y t x$$

Second, the transition function is required to return non-empty M structures.

3 Specification of the Stylized Decision Process

We specify the stylized GHG emissions decision process of the introduction in the theory summarized in Section 2. As a first step, we have to define the uncertainty monad M . Our decision process is a *stochastic* process and thus

$$M = \text{SimpleProb}$$

Here, *SimpleProb* is a finite probability monad: for an arbitrary type A , a value of type *SimpleProb* A is a list of pairs (A, Double_+) together with a proof that the sum of the *Double*₊ elements of the pairs is positive. These are double precision floating point numbers with the additional restriction (remember Section A.7) of being non-negative.

3.1 States, Controls

Second, we have to specify the states of the decision process. Consistently with Section 1.3 and with the notation introduced in Appendix we define:

$$\begin{aligned} \text{data } \text{State} = \text{DHU} | \text{DHC} | \text{DLU} | \text{DLC} | \text{SHU} | \text{SHC} | \text{SLU} | \text{SLC} \\ X t = \text{State} \end{aligned}$$

Third, we have to specify the controls of the decision process. In the introduction, we said that in states in which a green transition has not already been started (that is, in D -states), the decision maker has the option of either starting or further delaying the transition⁴

$$\text{data } \text{StartDelay} = \text{Start} | \text{Delay}$$

However, if a green transition has already been started, the decision maker has no alternatives. We formalize this idea by defining the set of controls in S -states to be a singleton. It will be useful to have two functions that test if a state is committed to impacts from climate change and if the economic wealth has taken a downturn:

⁴ Notice that we are using the term *transition* to denote two different notions: the green transition of the decision process and the function *next* of the theory discussed in Section 2 in which we now specify such a process!

$isCommitted, isDisrupted : (t : \mathbb{N}) \rightarrow X t \rightarrow Bool$

The idea is that $isCommitted$ ($isDisrupted$) returns *True* in C -states (L -states) and *False* in U -states (H -states).

3.2 The Transition Function

Finally, we have to specify the transition function of the process. As discussed in the introduction, this is defined in terms of transition probabilities.

The probabilities of starting a green transition Let’s first specify the probability that a green transition is started, conditional to the decision taken by the decision maker. Let

$$p_{S|Start} : Double_+$$

denote the probability that a green transition is started (during the time interval between the current and the next decision step) given that the decision maker has decided to start it. For a perfectly effective decision maker, $p_{S|Start}$ would be one. Let’s assume a 10% chance that a decision to start a green transition fails to be implemented, perhaps because of inertia of legislations, as discussed in Section 1.3:

$$p_{S|Start} = 0.9$$

Consistently, the probability that a green transition is delayed even if the decision maker has chosen to start it is

$$p_{D|Start} : Double_+$$

$$p_{D|Start} = 1.0 - p_{S|Start}$$

Similarly, we denote with $p_{D|Delay}$ and $p_{S|Delay}$ the probabilities that a green transition is delayed (started) given that the decision maker has decided to delay it. As a first step, we take $p_{S|Delay}$ to be equal to $p_{D|Start}$

$$p_{D|Delay} : Double_+$$

$$p_{D|Delay} = 0.9$$

$$p_{S|Delay} : Double_+$$

$$p_{S|Delay} = 1.0 - p_{D|Delay}$$

but we will come back to this choice in Section 6.

The probabilities of economic downturns In the informal description of the decision process from Section 1.3, we said that an essential trait of the decision process is that

...decisions to start a green transition, if implemented, are more likely to yield states with a low level of economic wealth (L -states) than states with high eco-

nomical wealth. This assumption reflects the fact that starting a green transition requires more investments and costs than just moving to states in which most of the work towards a globally decarbonized society remains to be done.

We need to formulate this idea in terms of transition probabilities. Let $p_{L|S,DH}$ denote the probability of transitions to states with a low level of economic wealth (L) given that a green transition has been started (S) from delayed states (D) with a high level of economic wealth (H). Similar interpretations hold for $p_{L|S,DL}$, $p_{L|S,SH}$, $p_{L|S,SL}$ and their counterparts for the cases in which a green transition has been delayed, $p_{L|D,DH}$ and $p_{L|D,DL}$. Remember that in our decision process

...once a transition has been started, it cannot be halted or reversed by later decisions or events.

In terms of transition probabilities, this means that we do not need to specify $p_{L|D,SH}$ and $p_{L|D,SL}$ because the probability of transitions from S -states to D -states is zero. We encode the requirement that “decisions to start a green transition, if implemented, are more likely to yield states with a low level of economic wealth (L -states) than states with high economic wealth” by the specification

$$pSpec3 : p_{H|S,DH} \leq p_{L|S,DH}$$

Because $p_{H|S,DH} = 1 - p_{L|S,DH}$, this requires $p_{L|S,DH}$ to be greater or equal to 50%. Let’s say that

$$p_{L|S,DH} = 0.7$$

We also want to express the idea that starting a green transition in a weak economy (perhaps a suboptimal decision?) is more likely to yield a weak economy than starting a green transition in a strong economy

$$pSpec4 : p_{L|S,DH} \leq p_{L|S,DL}$$

which requires specifying a value of $p_{L|S,DL}$ between 0.7 and 1.0, say

$$p_{L|S,DL} = 0.9$$

This fixes the values of $p_{L|S,DH}$ and $p_{L|S,DL}$ for our decision process in the ranges imposed by the “semantic” constraints $pSpec3$ and $pSpec4$. We discuss how these (and other) transition probabilities would have to be estimated in a more realistic (as opposed to stylized) GHG emissions decision process in Section 3.3.

Next, we have to specify the remaining transition probabilities $p_{L|S,SH}$, $p_{L|S,SL}$, $p_{L|D,DH}$ and $p_{L|D,DL}$. What are meaningful constraints for these? Remember that $p_{L|S,SH}$ and $p_{L|S,SL}$ represent the probabilities of transitions to low wealth

states (L -states) from H and L -states, respectively, while an already started green transition is accomplished. In this situation, and again because of the inertia of economic systems, it is reasonable to assume that transitions from H -states (booming economy) to H -states are more likely than transitions from H -states to L -states and, of course, the other way round. In formulas:

$$pSpec5 : p_{L|S,SH} \leq p_{H|S,SH}$$

$$pSpec6 : p_{H|S,SL} \leq p_{L|S,SL}$$

Again, because $p_{H|S,SH} = 1 - p_{L|S,SH}$ (and $p_{H|S,SL} = 1 - p_{L|S,SL}$), this requires $p_{L|S,SH}$ and $p_{L|S,SL}$ to be below and above 50%, respectively.

In our decision process, a high value of $p_{L|S,SL}$ implies a low probability of recovering from economic downturns in states in which a transition towards a globally decarbonized society has been started or has been accomplished. In more realistic specifications of GHG emission processes, one may want to distinguish between these two cases, or even to keep track of the time elapsed since a green transition was started and define the probability of recovering from economic downturns accordingly⁵.

Conversely, a low value of $p_{L|S,SH}$ means high *resilience* against economic downturns in states in which a transition towards a globally decarbonized society has been started or has been accomplished. In such states, we assume a moderate likelihood of fast recovering from economic downturns:

$$p_{L|S,SL} = 0.7$$

and also a moderate resilience

$$p_{L|S,SH} = 0.3$$

Let's turn the attention to the last two transition probabilities that need to be specified in order to complete the description of the transitions leading to economic downturns or recoveries. These are $p_{L|D,DH}$ and $p_{L|D,DL}$.

The semantics of $p_{L|D,DH}$ and $p_{L|D,DL}$ should meanwhile be clear: $p_{L|D,DH}$ represents the probability of economic downturns and $1 - p_{L|D,DL}$ the probability of recovering (from economic downturns) in states in which a green transition has not already been started. As for their counterparts discussed above, we have the semantic requirements

$$pSpec7 : p_{L|D,DH} \leq p_{H|D,DH}$$

$$pSpec8 : p_{H|D,DL} \leq p_{L|D,DL}$$

with $p_{H|D,DH} = 1 - p_{L|D,DH}$ and $p_{H|D,DL} = 1 - p_{L|D,DL}$ and thus, by the same argument as for $p_{L|S,SH}$ and $p_{L|S,SL}$, $p_{L|D,DH}$ and $p_{L|D,DL}$ below and above 50%, respectively.

How should $p_{L|D,DH}$ and $p_{L|D,DL}$ compare to $p_{L|S,SH}$ and $p_{L|S,SL}$? Is the likelihood of economic downturns in states in which a green transition has not already been started higher or lower than the likelihood of economic downturns in states in which a transition towards a globally decarbonized society has been started or has been accomplished? Realistic answers to this question are likely to depend on the decision step and on the time elapsed since the green transition has been started, see Section 3.3. As a first approximation, here we just assume that these probabilities are the same:

$$p_{L|D,DL} = p_{L|S,SL}$$

$$p_{L|D,DH} = p_{L|S,SH}$$

This completes the discussion of the probabilities of economic downturns and recoveries.

The probabilities of commitment to severe impacts from climate change

The last ingredient that we need to fully specify the transition function of our decision process are the probabilities of transitions to states that are committed to severe impacts from climate change. In the introduction, we have stipulated that

...we assume that the probability of entering states in which the world is committed to future severe impacts from climate change is higher in states in which a green transition has not already been started as compared to states in which a green transition has been started.

We account for this assumption with four transition probabilities: $p_{U|S,0}$, $p_{U|D,0}$, $p_{U|S}$ and $p_{U|D}$. The first two represent the probabilities of transitions (from uncommitted states) to uncommitted states at decision step zero for the cases in which a transition to a decarbonized economy has been implemented and delayed, respectively. Similarly, $p_{U|S}$ and $p_{U|D}$ represent the probabilities of transitions from U -states to U -states at later decision steps. We take the informal specification from Section 1.3 of the introduction

...delaying transitions to decarbonized economies increases the likelihood of entering states in which the world is committed to future severe impacts from climate change.

by the letter and, for the sake of simplicity, assume that the whole increase in the likelihood of entering committed states takes place during the first step of our decision process. This is a very crude assumption and we will come back to it when we discuss the results of measures of responsibility

⁵ As explained in the introduction, the main purpose of this paper is to present a novel approach towards measuring responsibility when decisions are to be taken under uncertainty. To this end, considering more realistic emission processes would be an unnecessary distraction.

in Section 4.4. With these premises (and keeping in mind that $p_{C|S,0} = 1 - p_{U|S,0}$, $p_{C|D,0} = 1 - p_{U|D,0}$, etc.) our informal specification translates into the constraints:

- $pSpec9 : p_{C|S,0} \leq p_{U|S,0}$
- $pSpec10 : p_{C|S,0} \leq p_{C|D,0}$
- $pSpec11 : p_{C|S} \leq p_{U|S}$
- $pSpec12 : p_{C|S} \leq p_{C|D}$
- $pSpec13 : p_{C|D,0} \leq p_{C|D}$

For the time being, we set $p_{U|S,0}$, $p_{U|D,0}$, $p_{U|S}$ and $p_{U|D,0}$ to 0.9, 0.7, 0.9 and 0.3, respectively. In words, we assume a 30% chance of committing to future severe impacts from climate change if we fail to start a green transition at the first decision step. We assume this chance to increase to 70% at later decision steps. We also assume a 10% chance of severe climate change impacts if we start a green transition at the first decision step or later. We will come back to these numbers in Section 6.2.

The transition function With the transition probabilities in place, we can now specify the transition function of the decision process. We proceed by cases, starting from transitions at step zero. The first case is the one in which the initial state is *DHU* and the decision was to start a green transition:

$next\ Z\ DHU\ Start = mkSimpleProb$

- $[(DHU, p_{D|Start} * p_{H|D,DH} * p_{U|D,0}),$
- $(DHC, p_{D|Start} * p_{H|D,DH} * p_{C|D,0}),$
- $(DLU, p_{D|Start} * p_{L|D,DH} * p_{U|D,0}),$
- $(DLC, p_{D|Start} * p_{L|D,DH} * p_{C|D,0}),$
- $(SHU, p_{S|Start} * p_{H|S,DH} * p_{U|S,0}),$
- $(SHC, p_{S|Start} * p_{H|S,DH} * p_{C|S,0}),$
- $(SLU, p_{S|Start} * p_{L|S,DH} * p_{U|S,0}),$
- $(SLC, p_{S|Start} * p_{L|S,DH} * p_{C|S,0})]$

In the above definition, *mkSimpleProb* is a function that (for an arbitrary type *A*) takes a list of pairs (*A*, *Double*₊) and returns a value of type *MA = SimpleProb A* that is, a finite probability distribution on *A*. The sum of the probabilities of the list elements has to be strictly positive; thus, the resulting probability distributions are sound per construction.

The interpretation of *next Z DHU Start* is straightforward given the transition probabilities introduced in the previous paragraphs. We only comment the definition of the probability of *SHU*, the state in which a green transition has been started, the economy is in a wealthy state and the world is not committed to future severe impacts from climate change.

This probability is defined by the product of three transition probabilities: the probability that a green transition is actually implemented, given that the decision was to do so $p_{S|Start}$; the probability that the economy is in a good state (an *H*-state) given that a green transition has been started from an *H*-state $p_{H|S,DH}$; and the probability of entering states that are not committed to severe impacts from climate change, again given that a transition to a decarbonized economy has been started $p_{U|S,0}$.

Notice that $p_{C|D,0} + p_{U|D,0}$ and $p_{C|S,0} + p_{U|S,0}$ are equal to one by definition of $p_{C|D,0}$ and $p_{C|S,0}$. The same holds for $p_{H|D,DH} + p_{L|D,DH}$ and $p_{H|S,DH} + p_{L|S,DH}$ (by definition of $p_{H|D,DH}$, $p_{H|S,DH}$) and for $p_{D|Start} + p_{S|Start}$ (by definition of $p_{D|Start}$). It follows that the sum of the probabilities of *next Z DHU Start* is one, as one would expect.

We can derive the probability of *SHU* (and of all other possible next states) given the decision to *Start* a green transition in *DHU*:

$$p_{S|Start} * p_{H|S,DH} * p_{U|S,0}$$

rigorously if we represent our decision process as a Bayesian belief network. To this end, it is useful to introduce some notation from elementary probability theory. Different textbooks adopt slightly different notations; here, we follow [40] and denote the conditional probability of entering *SHU* given the decision to *Start* a green transition in *DHU* with $P(SHU | Start, DHU)$. Thus, our obligation is to show

$$P(SHU | Start, DHU) = p_{S|Start} * p_{H|S,DH} * p_{U|S,0}$$

Let x_1, x_2, x_3 denote the “components” of the *current* state $x : X\ t$ and x'_1, x'_2, x'_3 the components of the *next* state. Thus, for $x = DHU$, we have $x_1 = D$, $x_2 = H$ and $x_3 = U$. As usual, we denote a decision in x at step t with $y : Y\ t\ x$.

The *variables* $x_1, x_2, x_3, y, x'_1, x'_2, x'_3$ and the decision step t are associated with the *nodes* of the Bayesian network of Fig. 1. The *edges* of the network encode the notion of conditional dependency: the arrow between x_1 and x'_2 posits that the probability of transitions to states with a low (high) economic wealth depends on whether a green transition is currently underway or has been delayed⁶.

The conditional probability tables associated with the nodes encode such probabilities. Thus, for example, the table associated with x'_1 posits that the conditional probability of entering *S*-states given that the decision (variable y) was to *Start* a green transition is $p_{S|Start}$ as discussed above. Similarly, the table associated with x'_2 encodes the specification

⁶ Because of the arrows from x_2 and x'_1 to x'_2 , such probability also depends on whether the current state of the economy is low or high and on whether a green transition gets started or not.

that the probability of entering an L -state given that an S -state was entered from a current D - and H -state is $p_{L|S,DH}$ ⁷. We can now derive $P(SHU|Start, DHU)$ from the Bayesian network representation of our decision process by equational reasoning. The computation is straightforward but we spell out each single step for clarity:

$$\begin{aligned}
 &P(SHU|Start, DHU) \\
 &= \text{-- definition of } x'_1 \dots y \dots x_3 \\
 &P(x'_1 = S, x'_2 = H, x'_3 = U \mid y = Start, x_1 = D, x_2 = H, x_3 = U) \\
 &= \text{-- definition of conditional probability, set theory} \\
 &P(x'_2 = H, x'_3 = U, x'_1 = S \mid y = Start, x_1 = D, x_2 = H, x_3 = U) \\
 &= \text{-- chain rule} \\
 &P(x'_2 = H \mid x'_3 = U, x'_1 = S, y = Start, x_1 = D, x_2 = H, x_3 = U) * \\
 &P(x'_3 = U, x'_1 = S \mid y = Start, x_1 = D, x_2 = H, x_3 = U) \\
 &= \text{-- chain rule} \\
 &P(x'_2 = H \mid x'_3 = U, x'_1 = S, y = Start, x_1 = D, x_2 = H, x_3 = U) * \\
 &P(x'_3 = U \mid x'_1 = S, y = Start, x_1 = D, x_2 = H, x_3 = U) * \\
 &P(x'_1 = S \mid y = Start, x_1 = D, x_2 = H, x_3 = U) \\
 &= \text{-- Bayesian network (conditional independence)} \\
 &P(x'_2 = H \mid x'_1 = S, x_1 = D, x_2 = H) * \\
 &P(x'_3 = U \mid x'_1 = S, x_3 = U) * \\
 &P(x'_1 = S \mid y = Start) \\
 &= \text{-- Bayesian network (tables)} \\
 &P_{H|S,DH} * P_{U|S,0} * P_{S|Start}
 \end{aligned}$$

Similar derivations can be obtained, in terms of the network of Fig. 1, for the other transition probabilities that define *next Z DHU Start* and, in fact, for all the transition probabilities that define *next*. Thus, Fig. 1 is in fact a compact representation of the transition function *next* of our decision process. Notice that the causal networks at the core of the storyline approach [10] are also Bayesian belief networks, albeit without a clearcut distinction between state and control spaces.

The case in which the initial state is *DHU* and the decision was to delay a green transition is similar to the first case with $p_{D|Start}$ and $p_{S|Start}$ replaced by $p_{D|Delay}$ and $p_{S|Delay}$, respectively:

⁷ Notice that the conditional probability table associated with x'_2 contains an undefined value α . This is because the probability of entering L (or H) states given that a D -state was entered starting from an S -state is irrelevant: the probability of transitions from S -states to D -states is zero (remember that we have assumed that green transitions cannot be halted or reversed by later decisions), as encoded in the third row of the table associated with x'_1 .

$$\begin{aligned}
 &next Z DHU Delay = mkSimpleProb \\
 &[(DHU, p_{D|Delay} * p_{H|D,DH} * p_{U|D,0}), \\
 &(DHC, p_{D|Delay} * p_{H|D,DH} * p_{C|D,0}), \\
 &(DLU, p_{D|Delay} * p_{L|D,DH} * p_{U|D,0}), \\
 &(DLC, p_{D|Delay} * p_{L|D,DH} * p_{C|D,0}), \\
 &(SHU, p_{S|Delay} * p_{H|S,DH} * p_{U|S,0}), \\
 &(SHC, p_{S|Delay} * p_{H|S,DH} * p_{C|S,0}), \\
 &(SLU, p_{S|Delay} * p_{L|S,DH} * p_{U|S,0}), \\
 &(SLC, p_{S|Delay} * p_{L|S,DH} * p_{C|S,0})]
 \end{aligned}$$

The cases in which the initial states are *DHC*, *DLU*, *DLC*, *SHU*, *SHC*, *SLU* and *SLC* are analogous to the *DHU* case and complete the specification of the transition function at decision step zero. The transition function at step one or greater is perfectly analogous with $p_{U|D}$, $p_{C|D}$, $p_{U|S}$ and $p_{C|S}$ in place of $p_{U|D,0}$, $p_{C|D,0}$, $p_{U|S,0}$ and $p_{C|S,0}$, respectively. Interested readers can find the full specification of the transition function [41], see file “Specification.lidr” in folder “2021. Responsibility under uncertainty: which climate decisions matter most?”

3.3 Realistic and Stylized Decision Processes


Before defining how much decisions under uncertainty matter in the next section, let us clarify the notion of *stylized* decision process. As mentioned in the introduction, this notion was originally introduced in [32] to contrast the one of *realistic* decision process. This is also the sense in which it has been used in this work.

For example, in discussing the probability of economic downturns, we have argued that, in the specification of more realistic GHG emissions decision processes, one might want to distinguish between states in which a transition towards a globally decarbonized society is ongoing and states in which the transition has already been accomplished.

In the case of ongoing green transitions, one may want to consider different transition probabilities, perhaps depending on the degree to which the transition has been accomplished or the time since it was started.

From this angle, more realistic essentially means a larger number of states (remember that, as discussed in the introduction, the states of a decision process typically represent sets of micro-states with the latter being detailed descriptions of physical, economic and social conditions), perhaps also of control options (for example, fast or slow green transitions) and hence more complex transition functions.

Fig. 1 Stylized decision process as a Bayesian network



This reductionist approach towards “realism” is paradigmatic of so-called *modelling* approaches. In climate policy advice, it has led to (integrated assessment) models of decision processes based on high-dimensional state and control spaces and a large number of model parameters [4, 11].

While this is popular in climate policy assessment and advice, the usage of “realistic” integrated assessment models (IAM) has also been criticized, among others, because of their poor understandability and limited predictive capability. For example, in [14], it was found that very different estimates of the “right” social cost of carbon can be “obtained” by setting the values of certain IAM parameters (for example, discount factors and climate sensitivities) to specific, arbitrary but “plausible” values and Pindyck even argued that

IAM-based analyses of climate policy create a perception of knowledge and precision that is illusory and can fool policymakers into thinking that the forecasts the models generate have some kind of scientific legitimacy [14].

Similar concerns and the problem that a too strong focus on *reliability* may be unsuitable for climate decision making at regional scales, have been discussed in [10].

Another weakness of IAMs for climate policy is their strong bias towards deterministic modelling. With very few exceptions, these models assume that decisions (e.g., of starting a global green transition) are implemented with certainty, that crucial parameterizations of climate processes (like the *equilibrium climate sensitivity*) can be estimated accurately and that the costs and the benefits of future

climate changes can be accounted for in suitable “terminal” (salvage, scrap, see [16] section 2.1.3) rewards.

Is there a way of specifying decision processes that are useful for pragmatic climate decision making and that avoid the drawbacks of deterministic modelling approaches based on high-dimensional state spaces?

We believe that this is the case and that, rather than neglecting uncertainty, the way to address this challenge is to 0) specify low-dimensional state and control spaces that are logically consistent with the informal description of the specific decision process at stake; 1) explicitly account for the uncertainties that are known to affect best decisions for that process, 2) exploit the knowledge available (from past experience, data, model simulation, etc.) to specify trustable transition probabilities with interpretations that are consistent with that process.

This is the essence of the approach that we have demonstrated in this section: starting from the informal description of Section 1.3, we have introduced formal specifications of state and control spaces that are logically consistent with that description. We have accounted for all the uncertainties of the informal description in terms of 12 transition probability parameters. For each parameter, we have provided an *interpretation* together with a range of values compatible with that interpretation. Within these ranges, we have then chosen certain values and defined the transition function in terms of those values. For example, we have postulated a 10% chance that a decision to start a green transition fails to be implemented.

In a (more) realistic specification, this figure could perhaps have been obtained by asking a pool of experts,

perhaps political scientists, historians, etc. Similarly, in more realistic specifications, the probabilities of recovering from economic downturns might be obtained from climate economists. These, in turn, might rely on model simulations, expert elicitation or perhaps statistical data. Finally, climate models (general circulation models, intermediate complexity models, low-dimensional systems of ordinary differential equations representing global mass and energy budgets) might be applied to representative micro-states samples of a given (macro) state (for example, our initial state *DHU*) to compute more realistic estimates (for example via Monte Carlo simulations) of transition probabilities, for instance, to committed states.

From this angle, the approach of “stylized” decision processes is similar to the *storyline* approach — the “identification of physically self-consistent, plausible pathways” — proposed in [10]. The focus, there on physical consistency and causal networks, is here on logical consistency and decision networks. Common to both approaches is the need to integrate contributions from very different disciplines, ranging from theoretical computer science to the social sciences [10, 42].

In this enterprise, the theory of Section 2 and the language extensions to be discussed in Sections 4 and 5 play a twofold role. On the one hand, they help ensure that results of model simulations, expert opinions, and statistical data are applied consistently. On the other hand, they make it possible to reason about pragmatic decision processes in a formal and rigorous way. We demonstrate this second aspect in Section 4.

4 Responsibility Measures

We formulate and answer three questions that we raised, informally, in the introduction:

- What does it mean precisely for decisions to matter?
- Are there general ways to measure how much decisions matter when these have to be taken under uncertainty?
- Is there a natural way of comparing similar decisions at different times?

We extend the theory of Section 2 with a *responsibility* measure for sequential decision processes under monadic uncertainty. The measure is obtained, for a given decision process, in three steps.

- S1 First, we need to define *the goal for which* we seek to measure responsibility, e.g., “saving the child” or “avoiding states that are committed to severe climate

change impacts”. We do this by extending our decision process to a full-fledged decision problem (compare Section 2).

- S2 Verified “best” and “conditional worst” decisions are compared at the specific state at which we want to measure how much decisions matter for the goal encoded in S1.
- S3 We define a degree of responsibility consistent with this measure.

This is how the theory of Section 2 is applied to implement the idea outlined in Section 1.2 for measuring how much decisions under uncertainty matter. For concreteness, we illustrate S1–S3 for the decision problem of Section 3. The extensions of the theory discussed in this section, however, are fully generic and can be applied to arbitrary decision processes. First, however, let’s discuss a general condition any measure of how much decisions matter (for whatever goal) should satisfy.

4.1 When Decisions Shall Not Matter

A responsibility measure has to attribute a non-negative number to the states of a decision process (e.g., the GHG emissions decision process specified in Section 3):

$$mMeas : (t : \mathbb{N}) \rightarrow X t \rightarrow Double_+$$

The idea is that $mMeas t x$ represents how much decisions in state x (at step t) do matter: the larger, the more the decisions in x matter. For the time being, assume that $mMeas t x$ takes values between zero and one. Under which conditions shall we require it to be zero? Certainly, we would like $mMeas t x$ to be zero whenever only one option is available to the decision maker in x :

$$mMeasSpec1 : (t : \mathbb{N}) \rightarrow (x : X t) \rightarrow Singleton (Y t x) \rightarrow mMeas t x = zero$$

Here, we have formalized the condition that only one option is available to the decision maker in x with the predicate *Singleton* ($Y t x$). We do not need to be concerned with the exact definition of *Singleton*: it is a component of our language and *Singleton A* posits that there is only one value of type A in a concise and precise way.

The specification $mMeasSpec1$ is consistent with *avoidance*, one of the three conditions put forwards in [33] under which “a person can be ascribed responsibility for a given outcome”. The other two conditions are *agency* (the capability to act intentionally, to plan, and to distinguish between desirable and undesirable outcomes) and *causal relevance*.

The notion of causality is not uncontroversial [43] and its role in formalizations of responsibility has been addressed, among others by [24, 25] and [26]. In the next section we show that, at least for sequential decision processes, it is possible to define “meaningful” measures of how much decisions matter without having to deal with causality. In Section 4.4, we discuss the relation between these measures and responsibility measures.

4.2 S1: Encoding Goals of Decision Making

To measure how much decisions matter *with respect to a specific goal*, we extend our decision process to a decision problem by encoding this goal into definitions for the components *Val*, *reward*, *meas*, \oplus , \leq and *zero* discussed in Section 2. The theory then allows the computation of *best* (and, as will be described below, *conditionally worst*) decisions for attaining *that given goal*. From these we will then define the measure of responsibility.

Note, however, that the decision problem thus specified is just a means to enable the definition of our measure of responsibility, and does *not* depend on the aims and preferences of an actual decision maker.

For example, in our stylized GHG emissions decision process, we might be interested in measuring how much decisions matter for avoiding states that are committed to severe impacts from climate change. Or perhaps we want to measure how much decisions matter for avoiding climate change impacts but also economic downturns. This can be done by defining

$$reward\ t\ x\ y\ x' = \mathbf{if\ } isCommitted\ (S\ t)\ x' \mathbf{\ then\ } 0 \mathbf{\ else\ } 1$$

or

$$reward\ t\ x\ y\ x' = \mathbf{if\ } isCommitted\ (S\ t)\ x' \vee isDisrupted\ (S\ t)\ x' \mathbf{\ then\ } 0 \mathbf{\ else\ } 1$$

with *Val* = *Double*₊ and \oplus , \leq and *zero* set to their canonical values for non-negative double precision floating point numbers. A special attention has to be taken in defining the measure function *meas*. Here, we follow standard decision theory and take *meas* to be the expected value measure

$$meas = expectedValue$$

but see Section 5.3 for alternative formulations.

In Section 5, we discuss generic goal functions and show how to automate the definition of *Val*, *reward*, etc. for such functions.

4.3 S2: Measuring How Much Decisions Matter

With a goal (avoiding climate change impacts but also economic downturns) encoded via the reward function, we have now extended the decision process of Section 3 to a decision problem. This allows us to tackle the problem of measuring how much decisions in a state do matter for that goal. For concreteness, let’s consider the initial state *DHU* of our decision problem. In this state, the decision maker has two options: start a green transition or further delay it. Remember that our decision maker is effective only to a certain extent. As shown in Fig. 1, a decision to start a green transition may well yield a next state in which the transition has been delayed. According to Section 3, the probability of this event is $p_{D|Start}$, that is, 10%.

What does this uncertainty imply for the decision to be taken in the initial state *DHU*? Answering this question rigorously requires fixing a decision *horizon*. This is the number of decision steps of our decision process that we look ahead in order to measure how much decisions matter. Remember from Section 2 that the value of taking zero decision steps is always *zero* : *Val*, a problem-specific reference value that holds for every decision step and state at that step. Thus, if we look forward zero steps, no decision matters, independently of the decision step and state. But, for a strictly positive number of decision steps, we can formulate and rigorously answer the following two questions

1. Is it better, in *DHU* to (decide to) start or to delay a green transition?
2. How much does this decision matter (for avoiding climate change impacts but also economic downturns)?

To do so, we first apply generic backward induction from Section 2 and compute an optimal sequence of policies *ps* over the horizon.

Remember that *bi* fulfills *biLemma*⁸. This means that no other policy sequence entails better decisions (again, for the goal of avoiding climate change impacts but also economic downturns) than *ps*. Thus, we can compute a best decision and the (expected) value (of the sum of the rewards associated with avoiding climate change impacts and economic downturns) over a horizon of *n* steps for arbitrary states:

⁸ If \leq , \oplus , *meas*, etc. fulfill the specifications from Section 2.2, see [44] for full machine-checked proofs.

```

best : (t, n : ℕ) → X t → String
best t Z x    = "The horizon must be greater than zero!"
best t (S m) x =
  let ps = bi (S t) m in
  let p  = bestExt ps in
  let b  = p x in
  let vb = val (p :: ps) x in "Horizon, best, value : " ++ show (S m) ++ ", " ++ show b ++ ", " ++ show vb

```

What is a best decision in *DHU* for a horizon of only one step?

```

* Responsibility > : exec best 0 1 DHU
Horizon, best, value : 1, Delay, 0.468

```

This is not very surprising: from the definition of *next 0 DHU Start* from Section 3, the probability of entering states that are either economically disrupted or committed to severe impacts from climate change is 0.708. Thus, the expected value of deciding to start a green transition is only $1 - 0.708 = 0.292$

By contrast, the expected value of deciding to delay a green transition is 0.468, as seen above. As it turns out, one has to look forward at least over three decision steps (or, in our interpretation, about three decades) for the decision to start a green transition to become a best decision in *DHU*. We can apply the computation

```

bests : (t : ℕ) → List ℕ → X t → IO ()
bests t Nil      x = putStrLn "done!"
bests t (n :: ns) x = do putStrLn (best t n x)
                    bests t ns x

```

to study how best decisions vary with the horizon. Again, for $x = DHU$ one obtains:

```

* Responsibility > : exec bests 0 [1..8] DHU
Horizon, best, value : 1, Delay, 0.468
Horizon, best, value : 2, Delay, 0.635454
Horizon, best, value : 3, Start, 0.940669612
Horizon, best, value : 4, Start, 1.250012318344
Horizon, best, value : 5, Start, 1.533635393558128
Horizon, best, value : 6, Start, 1.790773853744118
Horizon, best, value : 7, Start, 2.022874449805313

```

As anticipated, the decision to start a green transition at the first decision step becomes a best decision for horizons of three or more decisions. The other way round: our decision maker would have to be very myopic (or, equivalently very

much discount future benefits) to conclude that delaying a green transition is a best decision in *DHU*.

But how much does this decision actually matter? To answer this question, we need to compare a best decision in *DHU* for a given time horizon to a conditional worst decision. What does “conditional worst” mean in this context? Again, for concreteness, let’s for the moment fix the horizon to 7 decision steps.

What is the value (again, in terms of the sum of the rewards associated with avoiding climate change impacts and economic downturns) of deciding to delay a green transition in *DHU*? There are different ways of answering this question, but a canonical one⁹ is to consider the consequences of deciding to delay a green transition at the first decision step in *DHU* and take later decisions optimally. The approach is canonical because it corresponds to a well-established notion: that of stability with respect to local, not necessarily small, perturbations. In game theory, the notion is often called “trembling hands” (see [45], section 2.8) and was originally put forward by R. Selten in 1975 [46]. In our specific problem, it corresponds to considering the impact of a mistake (trembling hands) at the decision point at stake (*DHU*) under the assumption that future generations will act rationally to avoid negative impacts from climate change and economic downturns.

If we denote our optimal policy sequence for a horizon of 7 steps by *ps*, we can compute the consequences of deciding to delay at the first decision step in *DHU* and then take later decisions optimally by replacing the first policy of *ps* with one that recommends *Delay* in *DHU*:

```

ps : PolicySeq 0 7
ps = bi 0 7

ps' : PolicySeq 0 7
ps' = (setInTo (head ps) DHU Delay) :: tail ps

```

⁹ We discuss alternative approaches in Section 5.3.

The function *setInTo* in the definition of *ps'* is a higher-order primitive: it takes a function (in this case the first policy of *ps*), a value in its domain and one in its codomain, and returns a function of the same type that fulfills the specification

$$(setInTo\ f\ a\ b)\ a = b \wedge Not\ (a = a') \rightarrow (setInTo\ f\ a\ b)\ a' = f\ a'$$

for all *f*, *a*, *a'* and *b* of appropriate type. With *ps'*, we can compute the value of deciding to delay a green transition at the first decision step in *DHU*:

```
* Responsibility > : exec show [val ps DHU, val ps' DHU]
"[0.222874449805313, 1.672795254555656]"
```

The difference between the value of *ps* and the value of *ps'* in *DHU* then is a measure of *how much decisions in DHU matter for avoiding climate change impacts and economic downturns over a time horizon of 7 decision steps*: the bigger this difference, the more the decision matters.

4.4 S3: Responsibility Measures

We have argued that the difference between the value of *ps* and the value of *ps'* in *DHU*, is a measure of how much decisions in *DHU* matter for avoiding climate change impacts and economic downturns over a time horizon of 7 decision steps. This argument is justified because:

- We have defined optimal policy sequences to be policy sequences that avoid (as well as it gets) climate change impacts and economic downturns (S1).
- Over 7 decision steps, *ps* is a *verified optimal* policy sequence.
- The best decision in *DHU* is to start a green transition¹⁰:

```
* Responsibility > : exec show (head ps DHU)
"Start"
```

- *ps'* is a sequence of policies identical to *ps* except for recommending *Delay* instead of *Start* in *DHU* and for the first decision step:

```
* Responsibility > : exec show (head ps' DHU)
"Delay"
```

These facts are sufficient to guarantee that the difference between the value of *ps* and the value of *ps'* in *DHU* is actually the difference between the value (in terms of avoided climate change impacts and economic downturns over 7 decision steps) of the best and of the worst decisions that can be taken in *DHU*.

¹⁰ Remember from Section A.3 that *head ps x* = (*head ps*) *x* and that *head* is a function that returns the first element of a vector. Thus, *head ps* is a policy and we can apply it to a state to obtain a control.

The computation and the definitions of *ps* and *ps'* suggest a refinement and an implementation of the measure of how much decisions matter *mMeas* put forward in the beginning of this section. First, we want *mMeas* to depend on a time horizon *n*. Second, we want *mMeas* to return plain double precision floating point numbers

```
mMeas : (t : ℕ) → (n : ℕ) → X t → Double
mMeas t Z x = 0.0
mMeas t (S m) x = let ps = bi (S t) m in
                  let v = toDouble (val (bestExt ps :: ps) x) in
                  let v' = toDouble (val (worstExt ps :: ps) x) in
                  v - v'
```

Remember that, in Section 4.2, we have encoded the goal of avoiding severe climate change impacts and economic downturns for which we compute *mMeas* through a function

```
reward t x y : X (S t) → Double+
```

that returns 0 for next states that are committed to severe climate change impacts and economically disrupted and 1 otherwise. In this formulation, the value 1 is completely arbitrary: it could be replaced by any other positive number and perhaps discounted. This suggests that measures of how much decisions matter should be normalized

```
mMeas : (t : ℕ) → (n : ℕ) → X t → Double
mMeas t Z x = 0.0
mMeas t (S m) x = let ps = bi (S t) m in
                  let v = toDouble (val (bestExt ps :: ps) x) in
                  let v' = toDouble (val (worstExt ps :: ps) x) in
                  if v == 0 then 0 else (v - v') / v
```

Notice that, in states in which the control set is a singleton, any policy has to return the same control. In particular, the best extension and the worst extension of any policy sequence have to return the same control. Therefore, *mMeas* fulfills the avoidance condition from [33] discussed in S1 *per construction*. As a consequence, in *S*-states, the measure is identically zero, independently of the time horizon:

```
* Responsibility > : exec show (mMeas 0 4 SHU)
"0"
```

```
* Responsibility > : exec show (mMeas 0 6 SLC)
"0"
```

Notice also that *mMeas* can be applied to estimate how much decisions matter at later steps of a decision process. For example, we can assess that, for our decision process and under a fixed time horizon, decisions in *DHU* at decision step 0 matter less than decisions in *DHU* at later steps:

* *Responsibility* > : *exec show (mMeas 0 7 DHU)*
"0.1730602684132721"

* *Responsibility* > : *exec show (mMeas 1 7 DHU)*
"0.5673067719100584"

* *Responsibility* > : *exec show (mMeas 3 7 DHU)*
"0.5673067719100584"

This is not surprising given that the best decision, in *DHU* and for a time horizon of 7 decision steps, is to start a green transition and that, as stipulated in the introduction and specified in Section 3.2 through

$$pSpec9 : P_{C|D,0} \leq P_{C|D}$$

the probability of entering states in which the world is committed to future severe impacts from climate change is higher in states in which a green transition has not already been started as compared to states in which a green transition has been started.

4.5 Wrap-up

Through S1, S2 and S3, we have introduced a measure of how much decisions under uncertainty matter that fulfills the requirements for responsibility measures put forward in the introduction. It accounts for all the knowledge which is encoded in the specification of a decision process, it is independent of the aims of a (real or hypothetical) decision maker and it is *fair* in the sense that all decisions (decision makers) are measured in the same way.

Thus, we introduce *mMeas* as a first example of responsibility measure. In the next section, we generalize it by introducing a small DSL for the specification of goals of sequential decision processes under uncertainty and discuss alternative definitions.

5 Generic Goal Functions and Responsibility Measures

In the last section we have introduced a measure *mMeas* of how much decisions under uncertainty matter. We have constructed *mMeas* for the decision process of Section 3 in three steps and we have seen that, for this problem, *mMeas* fulfills the requirements for responsibility measures put forward in the introduction.

Specifically, in S1-S3, we have introduced an ad hoc definition of the reward function *reward* in terms of the (implicit) goal of avoiding *L*- (low economic wealth) and *C*- (committed) states and we have defined *mMeas* in terms of

the normalized difference between the value of two policy sequences.

In this section we generalize this construction: we drop the ad hoc definition of *reward* from Section 4 and introduce instead a small DSL to express goals explicitly. The DSL is implemented as an extension of the theory from Section 2 and consists of two artifacts: an abstract syntax and an interpretation function *eval*. The reward function is then defined generically in terms of the interpretation.

5.1 A Minimal DSL for Specifying Goals

Remember the definition of *reward* from S1 of Section 4:

```
reward t x y x' = if isCommitted (S t) x' ∨
                    isDisrupted (S t) x'
                    then 0.0
                    else 1.0
```

and that *reward t x y x'* represents the reward associated with reaching state *x'* when taking decision *y* in state *x* at decision step *t*, rewards are non-negative double-precision floating point numbers (*Val = Double₊*) and the rules for adding and comparing rewards are the canonical operations for this type.

In this formulation, the goal (avoiding states that are committed or that have a low level of economic wealth) for which we measure how much decisions matter is stated implicitly through the definition of *reward*.

Instead, we want to hide the definition of *reward*. In the theory of Section 2, this is the function that has to be specified to express the goal of decision making. Implementing *reward* could be challenging for domain experts with little computer science background. We want to give them a means to avoid the implementation and at the same time the opportunity of putting forward the goal of decision making transparently. This can be done through the definition:

```
goal = Avoid isCommitted && Avoid isDisrupted
```

Here *Avoid* is a function that maps Boolean predicates to goals. It is the fourth constructor of the abstract syntax

```
data Goal : Type where
  Exit   : Region → Goal
  Enter  : Region → Goal
  StayIn : Region → Goal
  Avoid  : Region → Goal
  (&&)   : Goal → Goal → Goal
  (||)   : Goal → Goal → Goal
  Not    : Goal → Goal
```

to specify goals for decision processes that are informed by notions of sustainable development or management [27, 30]: such goals are typically phrased in terms of a verb (*avoid*,

exit, enter, stay within, etc.) and of a *region* (predicate, subset of states) that encode notions of planetary boundaries or operational safety¹¹. In our formalization, such regions are encoded by

Region : Type

Region = (t : ℕ) → Subset (X t)

where *Subset A* is an alias for *A* → Bool. Notice the usage of the conjunction && in the specification of *goal*. Its semantics, like the semantics of the other constructors of the syntax, is given by the interpretation function

eval : Goal → (t : ℕ) → (x : X t) → Y t x → X (S t) → Bool

eval (Exit r) t x y x' = let t' = S t in elem t x (r t) ∧ ¬(elem t' x' (r t'))

eval (Enter r) t x y x' = let t' = S t in ¬(elem t x (r t)) ∧ elem t' x' (r t')

eval (StayIn r) t x y x' = let t' = S t in elem t' x' (r t')

eval (Avoid r) t x y x' = let t' = S t in ¬(elem t' x' (r t'))

eval (g && g') t x y x' = eval g t x y x' ∧ eval g' t x y x'

eval (g || g') t x y x' = eval g t x y x' ∨ eval g' t x y x'

eval (Not g) t x y x' = ¬(eval g t x y x')

While the definition of *eval* is almost straightforward¹², domain experts do not need to be concerned with it. They just apply the constructors of *Goal* to specify the goal of decision making like in the definition of *goal* given above. The goal for which we measure how much decisions matter is then fully transparent and the rewards are a straightforward function of *eval goal*:

reward t x y x' = if eval goal t x y x' then 1.0 else 0.0

5.2 Degrees of Commitment, Fuzzy Predicates

In more realistic (as opposed to stylized, see Section 3.3) GHG emissions decision processes, states are not necessarily either fully committed or fully uncommitted to severe impacts from climate change and decision makers are confronted with many degrees of commitment, possibly infinitely many.

A similar situation holds for other predicates on states, like being *vulnerable* (or *adapted*) to climate change or for measures of economic growth or welfare. This raises the question of how to specify the goals of decision making in decision processes in which predicates like *isCommitted* do not return Boolean values but, for example, values in [0, 1]. In this

¹¹ For example, in [30], a partitioning of the state space into a *sunny* region and its *dark* complement is the starting point for the construction of a hierarchy of regions: shelters, glades, lakes, trenches and abysses, see Fig. 1 at page 7.

¹² In this definition, *elem t* takes a state *x* : X t and a Boolean function on states (a subset of states) *s* : Subset (X t) and applies *s* to *x*.

situation, a partitioning of the state space into regions is not immediately available and the specification of goals requires an extension both of the syntax *Goal* for encoding goals and of the interpretation function *eval* associated with this syntax.

Discussing such extensions here would go well beyond the scope of this paper. However, the problem of developing a DSL for expressing the goals of decision making (and defining reward functions that are consistent with such goals) for realistic decision processes is a crucial step towards rationalizing decision making in climate policy advice and we plan to tackle this problem in an upcoming work.

5.3 Some Caveats

With *mMeas* defined as in Section 4 and with *goal* : Goal specified as above, one can recover the results for the decision process of Section 3. Before we turn back to this process in the last section, let us discuss a few aspects of the responsibility measures discussed so far.

One important trait of these measures is that they are obtained by extending the decision process for which one wants to measure how much decisions matter to a fully specified finite horizon sequential decision problem. In comparison to approaches like those proposed in [24, 25] and, more recently, [34], this approach has both advantages and disadvantages.

From the conceptual point of view, the major advantages are simplicity and straightforwardness: in contrast to models of causality like those put forward in the works mentioned above, finite horizon sequential decision problems are conceptually simple and well understood. Also, for finite horizon sequential decision problems, we can compute *verified* best and worst policies. This guarantees that the results obtained for a specific problem are a logical consequence of the assumptions made for that problem and not of programming errors or numerical errors. Because all the assumptions underlying a specific problem are put forward explicitly via specifications like

goal = Avoid isCommitted && Avoid isDisrupted,

the approach also guarantees high standards of transparency. Simplicity and straightforwardness are also the main drawbacks of our approach: we can only derive responsibility

measures for decision processes that can be naturally extended to finite horizon sequential decision problems.

This is the case for the stylized GHG emissions decision process discussed throughout our work and, indeed, for many interesting problems in climate policy because, as pointed out in [23]:

Climate policy decisions are necessarily sequential decisions over time under uncertainty, given the magnitude of uncertainty in both economic and scientific processes, the decades-to-centuries time scale of the phenomenon, and the ability to reduce uncertainty and revise decisions along the way.

But it is not immediately obvious how our approach could be applied to measure how much decisions matter in situations in which collective decisions emerge from a potentially large number of individual decisions, e.g., mediated through certain widely accepted mechanisms like majoritarian rules like in voting processes.

Another important aspect of the measures of responsibility proposed in this work is the comparison between verified best and what we called “conditional worst” decisions at the specific state at which we want to measure responsibility. Remember that, in the definition of $mMeas$, v and v' are $val(bestExt\ ps :: ps)\ x$ and $val(worstExt\ ps :: ps)\ x$, respectively. Here, $x : X\ t$ is a state at decision step t , ps is a verified optimal sequence of policies for taking n decisions starting from step $t + 1$ and $n + 1$ is the decision horizon.

Due to the definition of $bestExt$, generic backward induction and $biLemma$ from Section 2.2, $bestExt\ ps :: ps$ is an optimal policy sequence and $bestExt\ ps$ is an optimal policy (a function from states to controls) at decision step t . Similarly $worstExt\ ps$ is a policy that guarantees

$$val(worstExt\ ps :: ps)\ x \leq val(p :: ps)\ x$$

for all $x : X\ t$ and $p : Policy\ t$. In other words, we compare “best” decision (given by $bestExt\ ps$) and “worst” decision (given by $worstExt\ ps$) in x conditional to future decisions being best ones.

This is crucial because the difference between best and worst decisions (and hence our estimates of how much decisions matter) at a given step and in a given state would in general be different if we assumed that future decisions are not taken optimally.

In the context of our decision problem, for example, we would come up with a different measure of responsibility for “current” decisions if we assumed that future generations do not care about avoiding negative impacts from climate change or economic downturns or, equivalently, that they do care but do not act accordingly. If there are reasons to believe that this is the case, the verified optimal policy

sequence ps in the definition of $mMeas$ has to be replaced with one which is consistent with such a belief. For example, if we believe that the next generation will act more myopically (or more farsighted) than for a horizon of n decision steps, we have to compute ps accordingly. This can be done using the verified methods of the [15] theory.

Finally, we want to flag the role of the measure of uncertainty $meas$ from Section 2 in the definition of val and thus of v and v' . In all computations shown in this paper we have taken $meas$ to be the *expected value measure* but other measures of uncertainty are conceivable and we refer interested readers to [15, 19] and [35].

6 The Impact of Uncertainties on Responsibility Measures

In Sections 4 and 5 we have discussed a new method for assessing how much decisions under uncertainty matter in specific states and at specific decision steps of time-discrete decision processes.

We have introduced a small domain-specific language to encode the goal of decision making in terms of simple verb-predicate clauses and implemented a generic function $mMeas$ that fulfills the avoidance condition put forwards in [33] per construction. In this section, we show that, for the stylized decision process of Section 3, our measure of responsibility also fulfills the third condition discussed by [33]: $mMeas$ is zero for decisions that are *causally irrelevant*.

Further, we discuss how uncertainties affect how much decisions matter for that process. We argue that understanding how uncertainties affect the importance of decisions in (relatively) simple problems is a pre-condition for studying more realistic problems like, for example, those tackled in [4, 12]. As a first step, we study the impact of uncertainties about the capability of decision makers to actually implement decisions on $mMeas$. As in Section 4, we focus on values of $mMeas$ in DHU , the initial state of the decision process.

6.1 The Impact of Uncertainty About the Effectiveness of Decision Makers

The results discussed in Section 4 have been obtained for decision makers who are 90% effective: $p_{S|Start} = p_{D|Delay} = 0.9$. Specifically, we have seen that in DHU at decision step 0 and for a horizon of 7 decision steps $mMeas$ was about 0.173

```
*Application > : exec show (mMeas 0 7 DHU)
"0.1730602684132721"
```

and that in states with no alternatives, $mMeas$ is zero

Table 1 $p_{S|Start}$, best decisions and responsibility measures in *DHU* at decision steps 0 and 1 and for a horizon of 7 decision steps

$p_{S Start}$	best decisions	$mMeas\ 0\ 7\ DHU$	$mMeas\ 1\ 7\ DHU$
1.0	<i>Start</i>	0.155	0.581
0.9	<i>Start</i>	0.173	0.567
0.8	<i>Start</i>	0.187	0.551
0.7	<i>Start</i>	0.196	0.530
0.6	<i>Start</i>	0.199	0.504
0.5	<i>Start</i>	0.195	0.469
0.4	<i>Start</i>	0.181	0.420
0.3	<i>Start</i>	0.153	0.348
0.2	<i>Start</i>	0.100	0.230
0.1	<i>Start, Delay</i>	0.000	0.000
0.0	<i>Delay</i>	0.138	0.337

* *Application* > : *exec show (mMeas 0 7 SHU)*
"0"

and thus fulfills the avoidance condition mentioned above. What if decisions become *causally irrelevant*? Remember that the transition function of the decision process from Section 3 is completely defined through products of the 12 conditional probabilities that define the tables of the nodes of the belief network of Fig. 1, see Section 3.2. Of these conditional probabilities, only $p_{S|Start}$ and $p_{D|Delay}$ depend on the decision to start or to delay a green transition. This implies that, in our decision process, decisions become “causally irrelevant” when

$$p_{S|Start} = p_{S|Delay} \quad \wedge \quad p_{D|Start} = p_{D|Delay}$$

Because $p_{D|Start}$ is equal to $1 - p_{S|Start}$ and $p_{S|Delay}$ is equal to $1 - p_{D|Delay}$, this is equivalent to

$$p_{S|Start} = 1 - p_{D|Delay}$$

and we can test whether *mMeas* fulfills the causality relevance from [33] by replacing the definitions of $p_{S|Start}$ and $p_{D|Delay}$ in Section 3 with definitions that make decisions causally irrelevant. For example, setting

$$p_{S|Start} = 0.9$$

$$p_{D|Delay} = 0.1$$

yields

* *Application* > : *exec show (mMeas 0 7 DHU)*
"0"

¹³ The third one, the “capability to act intentionally, to plan, and to distinguish right and wrong and good and bad”, is a property of decision makers rather than a feature of decision processes. It is relevant for the attribution of blame, praise, sanctions or retributions to specific individuals but irrelevant for our work.

as one would expect. The same results obtain for $p_{S|Start} = p_{D|Delay} = 0.5$ and for all decision processes in which the sum of $p_{S|Start}$ and $p_{D|Delay}$ is one.

Having ascertained that our measure of responsibility fulfills two of the three *natural* conditions put forward in [33]¹³, we can turn the attention to the question of how uncertainties on the capability of decision makers to actually implement decisions affect measures of responsibility.

Let’s start by observing that it is not very realistic to assume that decision makers are equally effective in implementing the decision to *Start* and to *Delay* green transitions: delaying means following a minimal resistance, “business as usual” path. By contrast, implementing a global green transition requires a significant level of coordination and mutual trust between global players, not to mention huge economic investments and legislative efforts.

It follows that it makes sense to study the impact of uncertainty about the effectiveness of decision makers by fixing $p_{D|Delay}$, the probability that a green transition is delayed given that the decision was to delay it, to a relatively high value, say 0.9, and vary $p_{S|Start}$. What happens to our measure of responsibility when $p_{S|Start}$ decreases? We have seen that, for $p_{S|Start} = p_{D|Delay} = 0.9$, the measure of responsibility for a horizon of 7 decision steps was about 0.173 in *DHU* and at decision step 0.

We know that *mMeas 0 7 DHU* has to become zero as $p_{S|Start}$ goes down to 0.1 ($p_{D|Delay}$ is fixed to 0.9) because for these values decisions become causally irrelevant. Does the measure of responsibility *mMeas 0 7 SHU* linearly decrease from 0.173 to 0 as $p_{S|Start}$ decreases from 0.9 to 0.1? Table 1 shows that, in contrast to the popular intuition that “if decisions can hardly become true, they do not matter after all”, this is not the case:

Far from being linear, the measure of responsibility is not even monotonous! For the case in which a decision to start a green transition is implemented with only 50% of probability, *mMeas 0 7 DHU* is actually higher than for the case in which such decision is realized with certainty. Computations of *mMeas 1 7 DHU* confirm these observations. In this case the responsibility decreases monotonically with $p_{S|Start}$ but, again, non-linearly.

Notice also that the best decision for $p_{S|Start} = 0$ is *Delay*. This is not surprising: the decision to delay a green transition implies a 10% probability that the transition is actually started. This is low but higher than 0, the probability that a green transition gets started if the decision was *Start*.

This concludes the study of the impact of $p_{S|Start}$ and $p_{D|Delay}$ on our measure of responsibility. Before turning the attention to the impact of uncertainties about commitment to severe impacts from climate change on *mMeas 0 7 DHU*, let’s remark that values of *mMeas 0 7* and *mMeas 1 7* in *DLU* are qualitatively similar to those in *DHU* albeit higher: for $p_{S|Start} = 0.2$, for example *mMeas 0 7 DLU* = 0.144, 44%

Table 2 Like Table 1 but with $p_{U|S,0}$, $p_{U|D,0}$, $p_{U|S}$ and $p_{U|D}$ set to 0.7, 0.5, 0.7 and 0.1 (instead of 0.9, 0.7, 0.9 and 0.3)

$p_{S Start}$	best decisions	$mMeas\ 0\ 7\ DHU$	$mMeas\ 1\ 7\ DHU$
1.0	Start	0.141	0.748
0.9	Start	0.159	0.733
0.8	Start	0.171	0.713
0.7	Start	0.177	0.689
0.6	Start	0.178	0.658
0.5	Start	0.170	0.616
0.4	Start	0.154	0.557
0.3	Start	0.125	0.468
0.2	Start	0.077	0.315
0.1	Start, Delay	0	0
0.0	Delay	0.098	0.485

higher than in initial states with high economic wealth. We do not show detailed results for $mMeas\ 0\ 7\ DLU$ and $mMeas\ 1\ 7\ DLU$ but these are available at [41]. We will come back to these observations in Section 7.

6.2 The Impact of Uncertainty About Commitment

In Section 3, we have accounted for the possibility of transitions to states that are committed to severe impacts from climate change in terms of four conditional probabilities $p_{U|S,0}$, $p_{U|D,0}$, $p_{U|S}$, $p_{U|D}$ and their complements.

Remember that $p_{U|S,0}$ represents the probability of entering uncommitted states right after the first decision step given that a green transition was implemented. Similarly, $p_{U|S}$ represents the probability of entering uncommitted states at later decision steps given that a green transition was implemented in those steps or earlier. Similarly for $p_{U|D,0}$ and $p_{U|D}$.

In all scenarios discussed so far $p_{U|S,0}$, $p_{U|D,0}$, $p_{U|S}$, $p_{U|D}$ were set to 0.9, 0.7, 0.9 and 0.3, respectively. This means assuming a 10% chance of committing to future severe impacts from climate change if we manage to start a green transition at the first decision step and a 30% chance if we fail to do so. We have also assumed that the chance of committing to future impacts from climate change if we fail to start a green transition increases from 30% at the first decision step to 70% at later decision steps.

This is perhaps a little bit too optimistic if we consider that, in the Oct. 2018 “Summary for Policymakers”, the [27] estimates that about 50% of the “pathways limiting global warming to 2 degrees Celsius with at least 66% probability” will attain zero net CO_2 emissions between about 2060 and 2080 whereas more ambitious paths (limiting global warming to 1.5 degrees Celsius) reach zero net CO_2 emissions earlier. The IPCC report suggests that a more realistic estimate of $p_{U|S,0}$ (if we identify our green transition corridor with

one that attains zero net CO_2 emissions between about 2060 and 2080 and associate commitment to severe impacts from climate change with violating the 2 degrees Celsius goal) would perhaps be about 0.66.

What if we assume $p_{U|S,0} = 0.7$ and lower $p_{U|D,0}$ accordingly, say to 50%? For consistency, we also need to decrease $p_{U|S}$ and $p_{U|D}$, say to 0.7 and 0.1. The corresponding measures of responsibility in DHU at decision steps zero and one and for an horizon of 7 steps are reported in Table 2.

By comparing these results with those of Table 1, we see that the effect of increasing the probability of severe impacts from climate change by 20% has been to systematically *decrease* how much decisions matter at the first decision step and to *increase* how much decisions matter at the second decision step. We will come back to this observation in the conclusion.

7 Conclusion

In this paper, we have studied the notion of responsibility under uncertainty in sequential decision processes in the context of global climate policy. Specifically, we have extended the verified theory of policy advice and avoidability [15] with a family of methods for measuring how much decisions under uncertainty do matter and the degree of responsibility associated with such decisions.

We have also introduced a small domain-specific language for specifying sustainability goals in GHG emissions decision processes. We have applied the DSL to formalize a stylized decision process in which a decision maker repeatedly faces two options over a finite number of decision steps: start a “green” transition to a decarbonized society or delay such transition. We have studied how uncertainties (on the capability of decision makers to actually implement their decisions and on the consequences of starting or delaying green transitions) affect how much decisions at specific points in time do matter and the degree of responsibility associated with these decisions.

Some of the results presented in Sections 4, 5 and 6 are consistent with common intuitions on how responsibility changes when the capability of decision makers to actually impose their decisions increases or decreases.

Perhaps more surprisingly, we have also found that the measures of responsibility discussed in Section 4 suggest that a “moral” approach towards decision making — doing the right thing even though the probability of success becomes increasingly small — is perfectly rational over a wide range of uncertainties.

The fact that these results are based on verified methods for computing optimal policies is crucial for their interpretation: they are a logical consequence of the assumptions about the decision process specified in Section 3 and of the

goals of decision making (avoiding short term economic downturns and long term negative impacts from climate change) explicitly stated in Section 5 and not the result of programming errors.

The fact that “best” decisions are stable with respect to both decision horizons (the number of decision steps to look forward in order to define measures of responsibility) and to the amount of uncertainty suggest that our results could be valid for more realistic decision processes than the one studied here.

In the last section, we have also shown that the measures of responsibility introduced in Section 4 fulfill two of the three natural conditions put forward in [33]. For the third condition, see footnote 14 on page 41.

Also in Section 6, we have shown that, in DHU (green transition delayed, economic welfare high, uncommitted to negative impacts from climate change) the importance of taking the right decision (starting a green transition) at decision step 0 systematically *decreases* (as compared to the importance of taking the right decision — also starting a green transition — at decision step 1) as the probability of severe impacts from climate change *increases*.

It is important to point out that this result is only in apparent contradiction with the intuition (that inspires, among others, the “Fridays for future” movement) that current climate decisions matter more than decisions to be taken in the upcoming decades. This is because of two reasons.

The first one is that the probability of facing the decision to either start or to delay a green transition in DHU at decision step 1 is less than one. In other words: it is true that, if the next generation will happen to be in DHU, they will face a decision that matters more than the current one. But the probability that the next generation *will be* in DHU is relatively low, especially if the current decision is to further delay a green transition!

The second reason why the results discussed in Section 6 are not in contradiction with the notion that current climate decisions matter more than decisions to be taken in the upcoming decades is more subtle and needs to be discussed with some care.

In the introduction, we have pointed out a fundamental difficulty of climate policy advice: the lack of agreement on how to account for the chances and the risks of climate change.

In the language introduced in Section 5, lack of agreement on how to account for the chances and the risks of climate change means lack of agreement on how to define *goal*. Remember that the results discussed in Section 6 have been obtained with

goal = Avoid isCommitted && Avoid isDisrupted

In other words, we have measured how much decisions matter and how to attribute responsibility to specific decisions

with respect to the goal of avoiding negative impacts from climate change and economic downturns.

This encodes notions of sustainability but not necessarily of *fairness* (balanced share of responsibility between generations), not to mention *justice*: there is nothing in the above definition of *goal* that prevents optimal decisions to lead to states in which the set of options available to upcoming generations has shrunk or to states in which decision makers have to face more crucial decisions than the current one.

By contrast, the idea that current climate decisions matter more than decisions to be taken in the upcoming decades is based on notions of fairness and justice that are not encoded in *goal* and thus are not accounted for in the analysis presented in Section 6. As far as one can define predicates on states that encode notions of fairness and justice, one can apply the measures of responsibility from Section 5.

The problem to agree on what is to be considered fair and just limits the applicability of rigorous decision theories to climate policy. But it is a problem that, to quote Hardin [47], “has no technical solution” and cannot be avoided — neither by verified decision making [15], nor by generalizations of cost-benefit analysis [42], multi-objective optimal control [13] or storylines [10].

From this perspective, this paper can also be seen as a contribution from verified decision theory towards understanding the limits of applicability of decision theories to policy advice.

Appendix. Functional Notation

A.1 Imperative, Functional and Dependently Typed Languages

Somewhat simplified, an *imperative* program is a sequence of instructions of what a computing machine should do. In contrast, a *functional* program is a description of what the machine should compute as a mathematical function from input to output. Common to both paradigms is the ability to name and reuse patterns of computation to enable concise and precise descriptions of algorithms.

Dependently typed languages like NuPRL, Coq, Agda, Idris or Lean support implementing programs but also postulating axioms, building theories and formulating program specifications. These are formal descriptions of what programs are required to do. Program specifications are crucial for verified programming. Verified programs are programs that have been machine-checked (verified) to fulfill a specification. They represent the highest correctness standard currently achievable [22, 49].

A.2 Expressions and Their Types

At the core of all programming languages is a sublanguage of *expressions* like $1 + 2$, "Hello", $[1, 7, 3, 8]$, etc. In strongly typed languages like Idris each "valid" expression has a *type*, like \mathbb{N} , *String*, *List* \mathbb{N} , etc. The judgment $e : t$ states that the expression e has type t . Most of the power of Idris comes from its type-checker which can check these judgments for very complex expressions e and types t . In the examples below we use a few arbitrary but fixed types A , B and C .

A.3 Function Application and Currying

In Idris (and several other functional languages like Haskell and Agda) the notation for function application is juxtaposition. You can think of it as an invisible infix operator binding more strongly than any other operator. Thus, $f x$ denotes the application of the function f to the argument x . Parentheses are used as in mathematics to resolve operator precedence like in $(2 + 3) * 4$ and to denote tuples like $(1, True, 'c')$. It is always possible to add extra parentheses, so $f (x)$ is also a valid syntax for function application.

In mathematics, a function of $n > 1$ arguments is often "implicitly converted" to a function taking as arguments n -tuples. For example, if g takes one argument in A and another one in B and returns values in C , we write $g (x, y)$ to denote the application of g to the pair $(x, y) : (A, B)$ (in Idris, (A, B) denotes the Cartesian product of A and B) and say that g has type $(A, B) \rightarrow C$.

In functional notation we instead use nested function application and write $(g x) y$ (which can also be written $g x y$ because function application is left-associative) to denote the application of g to $x : A$ and $y : B$. Thus, g has type $A \rightarrow (B \rightarrow C)$ or simply $A \rightarrow B \rightarrow C$, $g x$ has type $B \rightarrow C$ and $g x y : C$. This is called the *curried* form. Infix operators like $(+) : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}$ are, just as in mathematics, a special case where a (binary) function can be written between its first and second argument: $2 + 3 : \mathbb{N}$.

A.4 Definitions, Pattern Matching and Recursion

The ability to name and reuse expressions is at the core of all programming languages. In strongly typed functional languages, we can name and reuse expressions as long as we provide their type.

$aNumber : \mathbb{N}$

$aNumber = 1738$

Any time $aNumber$ is used we can just substitute 1738. We can define functions through lambda-expressions:

$aFun : \mathbb{N} \rightarrow \mathbb{N}$

$aFun = \lambda x \Rightarrow 2 * x + 1$

or, equivalently $aFun x = 2 * x + 1$. The latter form is useful when we want to distinguish different cases by *pattern matching*:

$(\uparrow) : Double \rightarrow \mathbb{N} \rightarrow Double$

$x \uparrow Z = 1$

$x \uparrow (S n) = x * (x \uparrow n)$

The two cases (for zero and the successor of n) can be seen as equations we want to hold for the "to the power of" binary operator (\uparrow) . In addition to pattern matching, this example also introduces recursion: the function being defined, (\uparrow) , is applied to $(S n)$ on the left hand side and to n on the right hand side of the second equation.

A.5 Partial Application and Higher-order Functions

If a function of two (or more) arguments, $g : A \rightarrow B \rightarrow C$, is applied to just one argument x we obtain a function $g x : B \rightarrow C$ which is a *partially applied* version of g . Thus, we can view any function as a 1-argument function, possibly returning a function.

We can also convert g into $h : (A, B) \rightarrow C$ by pairing up the first two arguments. More generally, we can convert any (n -argument) function into a 1-argument function that takes as arguments n -tuples. For binary functions this conversion can be done generically:

$uncurry : (A \rightarrow B \rightarrow C) \rightarrow ((A, B) \rightarrow C)$

$uncurry f (a, b) = f a b$

The implementation is straightforward: *uncurry* takes as input a function f which takes values of type A to functions from B to C . It returns a function that takes as input pairs of type (A, B) . This is our first example of a *higher-order* function: a function taking another function as a parameter. The opposite transformation is also short and clean:

$curry : ((A, B) \rightarrow C) \rightarrow (A \rightarrow B \rightarrow C)$

$curry f a b = f (a, b)$

These examples are a bit abstract, so here is a more applied example: Given a time-dependent reward function $reward : \mathbb{N} \rightarrow A \rightarrow Double$ and a parameter $rate : Double$ we construct a discounted reward function by applying the higher-order function *discount*:

$discount : Double \rightarrow (\mathbb{N} \rightarrow A \rightarrow Double) \rightarrow$

$(\mathbb{N} \rightarrow A \rightarrow Double)$

$discount rate reward = \lambda t \Rightarrow \lambda x \Rightarrow (rate \uparrow t) * (reward t x)$

A.6 Polymorphic Functions and Equality Types

The types presented so far have been *monomorphic*: using only specific types like \mathbb{N} , *Double* and the fixed types *A*, *B*, and *C*. Many programs work generically for a large class of types. For example, *discount* works for any *A* and *curry* for any *A*, *B*, and *C*. A simpler example is the projection function *fst* for extracting the first component of a pair:

```
fst : {A, B : Type} → (A, B) → A
fst (x, y) = x
```

The type of *fst* depends on two *type variables* *A* and *B*. Thus *fst* is in fact a three-argument function taking two types and a pair and returning the first component of the pair. The two first arguments are *implicit* arguments which can be inferred by the system in most use cases. In Section 2, most functions are polymorphic, using a combination of explicit and implicit type arguments.

Dependently typed functional languages support reasoning about the equality of expressions. The claim that an expression $a : A$ is equal to an expression $b : B$ is written simply $a = b$. The infix operator ($=$) has type $A \rightarrow B \rightarrow Type$ and defines a whole family of types: for every $a : A$, and $b : B$ we have a type $a = b$. Almost all types in this family are empty (uninhabited, contain no values) but a few contain one value written *Refl* : $a = a$. Thus, a value $p : a = b$ tells us that a and b are equal (and p is a proof of that fact). Here are two examples of using equality types to specify properties of multiplication:

```
multUnitSpec : (y : Double) → 1 * y = y
multAssocSpec : (x, y, z : Double) → (x * y) * z = x * (y * z)
```

A.7 Dependent Types and Data Declarations

Many programming languages use types to make sure the code doesn't go wrong, but dependently typed languages support types which depend on values. We have already seen some examples: *multUnitSpec* is a function whose return type $(1 * y = y)$ depends on y , a value of type *Double*. The equality type $x = y$ depends on the two values x and y .

We can take advantage of dependent types to specify requirements that multiplication shall fulfill, as in the *multUnitSpec* and *multAssocSpec* examples. We can also apply dependent types to restrict the values of arguments, for example, to specify a square root function that accepts only non-negative arguments. Restricting sounds negative, but it allows to avoid nonsensical combinations of values, which helps to eliminate whole classes of software bugs.

Let's start with a non-dependent data declaration that introduces natural numbers:

```
data N : Type where
  Z : N
  S : N → N
```

This states that \mathbb{N} is a type and that values of type \mathbb{N} can be constructed using *Z* for zero and *S n* for the successor of $n : \mathbb{N}$. The same declaration can be written in a less verbose form

```
data N = Z | S N
```

where the vertical bar separates the two *data constructors* *Z* and *S*. In this form, the types of \mathbb{N} , *Z* and *S* are implicit. We apply this form of data declaration in Section 3.1, e.g., to define the data types *State* and *StartDelay*.

With \mathbb{N} and the syntax for **data** declarations in place we move on to the more complex example of lists of fixed-length:

```
data Vect : N → Type → Type where
  Nil : Vect Z a
  (::) : (x : a) → (xs : Vect n a) → Vect (S n) a
```

Thus, for any $n : \mathbb{N}$ and $A : Type$, values of type *Vect n A* are lists of length n of elements of type *A*. For example

```
xs : Vect 3 Double
xs = 0.1 :: 0.6 :: 0.4 :: Nil
```

Idris also provides syntax extensions for defining vectors and lists of variable length in square brackets notation:

```
ys : Vect 4 N
ys = [1, 2, 3, 4]
```

```
zs : List String
zs = []
```

A simple example of a vector based function is *head* which extracts the first element of a vector:

```
head : {n : N} → {A : Type} → Vect (S n) A → A
head (x :: xs) = x
```

Note that *head* is only defined for non-empty vectors: vectors of length *S n* for some n . By restricting the arguments of *head*, we make sure that the function is never applied to empty vectors, thus eliminating a common source of errors.

In Section 2 we use a data declaration similar to *Vect* to define a datatype *PolicySeq* for sequences of policies of fixed-length.

A.8 Properties as Types, Specifications

In programming, the type *Bool* is often used to collect the two truth values *False* and *True* and to implement run-time “truth” tests. In dependently typed programming and constructive mathematics we can go one step further and represent truth “values” at the type level. Such truths can then be type-checked at compile time and before a (possibly incorrect) program is executed.

In much the same way as $(=) : A \rightarrow B \rightarrow Type$ represents equality of expressions at the type level (remember Section 1) and allows us to construct equality proofs straightforwardly,

$$p : 2 + 3 = 5$$

$$p = Refl$$

we can represent other binary relations through types. For example, we can define a “smaller or equal” relation for natural numbers

$$(\leq) : \mathbb{N} \rightarrow \mathbb{N} \rightarrow Type$$

or, as in Section 2.1, one that compares values in *Val*, the type used there to represent rewards. In all cases the idea is the same: any inhabited type (any type for which we can provide a value, like *Refl* for equality) represents truth and empty types represent falsity. An empty type, usually called *Void*, is easily defined through a **data** declaration with no constructors

data *Void* : *Type* **where**

and (\leq) can be defined in such a way that $7 \leq 3$ is empty (thus, no values of this type can be constructed) and values of type $3 \leq 7$ can be constructed easily, similarly to values of type $2 + 3 = 5$.

The function (\leq) is an example of a binary *predicate* on natural numbers. Similarly, unary predicates on values of type *A* can be represented by functions $P : A \rightarrow Type$. Predicates are useful to specify properties of computations, as we have already seen with *multUnitSpec* and *multAssocSpec*. Here is another example: given a sorting function

$$sort : \{n : \mathbb{N}\} \rightarrow \{A : Type\} \rightarrow Vect\ n\ A \rightarrow Vect\ n\ A$$

and a predicate representing “sortedness”:

$$Sorted : \{n : \mathbb{N}\} \rightarrow \{A : Type\} \rightarrow Vect\ n\ A \rightarrow Type$$

we can formulate the requirement that *sort* shall return sorted vectors:

$$sortSpec : \{n : \mathbb{N}\} \rightarrow (xs : Vect\ n\ A) \rightarrow Sorted\ (sort\ xs)$$

Table 3 Propositions-as-Types and Proofs-as-Programs (“Curry-Howard”) correspondence relating dependent type theory and logic [50, 51]

Functional notation (Idris)	Logic
$p : P$ (<i>p</i> is a program of type <i>P</i>)	<i>p</i> is a proof of <i>P</i>
inhabited type	provable proposition
empty type	False
singleton type	True
$P \rightarrow Q$	<i>P</i> implies <i>Q</i>
$Exists\ \{A\}\ P$	there exists a witness $x : A$ such that <i>P</i> <i>x</i> holds
$(x : A) \rightarrow P\ x$	for all <i>x</i> of type <i>A</i> , <i>P</i> <i>x</i> holds

Any valid implementation of *sortSpec* is then logically equivalent to a proof that, for any vector *xs*, *sort* *xs* is sorted¹⁴. More generally, any function of type $(x : A) \rightarrow P\ x$ (for any predicate $P : A \rightarrow Type$) is logically equivalent to a proof that *P* *x* is non-empty for every $x : A$. In Section 2 we exploit this equivalence to posit that the set of controls *Y* *t* *x* associated with a state $x : X\ t$ shall not be empty through the function *notEmptyY*.

A.9 Programs, Proofs, and Totality

We have seen that, in dependently typed languages, properties can be represented by types and proofs by values of these types. The correspondence between functional notation and logic goes deeper and we sum up the main results in Table 3:

When we embed logic in a dependently typed language, we have to require all our functions to be total (otherwise the logic will be inconsistent).

A total function $f : A \rightarrow B$ is defined for all $x : A$, whereas a partial function is undefined for some $x : A$. If partial functions were allowed, we could use them to prove any theorem, including patently false ones. A simple example is the partial function *headL* : *List* *A* → *A* which is undefined for empty lists. Using *headL* we could easily prove a false theorem (like $3 = 5$) by first building an empty list of proofs, $[] : List\ (3 = 5)$, and then extracting the first element:

$$surprise : 3 = 5$$

$$surprise = headL\ []$$

A function may cover all cases, but still fail to be total. An extreme example is the completely circular definition

¹⁴ An implementation of *sortSpec* is not enough to guarantee that *sort* is correct but is a first step in the right direction.

circular : Void

circular = *circular*

If we require functions to be total (which we do in Sections 2 to 6), the totality checker will warn about missing cases and potentially circular definitions.

Acknowledgements The authors thank the editors and reviewers, whose comments have led to significant improvements of the original manuscript. The work presented in this paper heavily relies on free software, among others on Coq, Idris, Agda, GHC, git, vi, Emacs, LaTeX and on the FreeBSD and Debian GNU/Linux operating systems. It is our pleasure to thank all developers of these excellent products.

Author Contributions NiBo drafted the initial sketch, performed the inlined computations and managed the preparation and the submission of the manuscript. NuBr drafted parts of Section 3 and implemented the DSLs of Section 5. MM and NuBr performed the sensitivity analysis of the responsibility measures. MC, CI and ZL revised the introduction, Appendix and Section 2 and the specification of the stylized GHG emissions decision problem. PJ drafted Appendix. TR improved the formulation of the sequential decision problems of Section 4. All authors read, commented and approved the final manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL. This is TiPES contribution no. 138. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 820970.

Availability of Data and Materials This manuscript has been written as literate Idris programs, one program per section. The programs/sections have been processed with lhs2TeX [48] and with latexmk to generate a PostScript manuscript. All files are publicly available at [41] in folder "2021.Responsibility under uncertainty: which climate decisions matter most?"

Code Availability At [41], see above.

Declarations

Ethics Approval Not applicable

Consent to Participate Not applicable

Consent for Publication Not applicable

Conflict of Interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Talbert, M. (2014). Moral responsibility. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2014 ed). Metaphysics Research Lab, Stanford University.
2. Pierrehumbert, R. T. (2010). *Principles of planetary climate*. Cambridge University Press, Cambridge, UK.
3. Lee, H., Muri, H., Ekici, A., Tjiputra, J., & Schwinger, J. (2021). The response of terrestrial ecosystem carbon cycling under different aerosol-based radiation management geoengineering. *Earth System Dynamics*, 12(1), 313–326. <https://doi.org/10.5194/esd-12-313-2021>
4. Helwegen, K. G., Wieners, C. E., Frank, J. E., & Dijkstra, H. A. (2019). Complementing CO₂ emission reduction by solar radiation management might strongly enhance future welfare. *Earth System Dynamics*, 10(3), 453–472. <https://doi.org/10.5194/esd-10-453-2019>
5. Moreno-Cruz, J., & Keith, D. (2012). Climate policy under uncertainty: A case for geoengineering. *Climatic Change*. <https://doi.org/10.1007/s10584-012-0487-4>
6. Eastham, S. D., Weisenstein, D. K., Keith, D. W., & Barrett, S. R. H. (2018). Quantifying the impact of sulfate geoengineering on mortality from air quality and UV-B exposure. *Atmospheric Environment*, 187, 424–434. <https://doi.org/10.1016/j.atmosenv.2018.05.047>
7. Zarnetske, P. L., Gurevitch, J., Franklin, J., Groffman, P. M., Harrison, C. S., Hellmann, J. J., Hoffman, F. M., Kothari, S., Robock, A., Tilmes, S., Vioni, D., Wu, J., Xia, L., & Yang, C.-E. (2021). Potential ecological impacts of climate intervention by reflecting sunlight to cool earth. *Proceedings of the National Academy of Sciences*, 118(15). Retrieved December 20, 2022, from <https://doi.org/10.1073/pnas.1921854118>. <https://www.pnas.org/content/118/15/e1921854118.full.pdf>
8. Bhowmick, M., Mishra, S. K., Kravitz, B., Sahany, S., & Salunke, P. (2021). Response of the Indian summer monsoon to global warming, solar geoengineering and its termination. *Scientific Reports*, 11(1), 1–11. <https://doi.org/10.1038/s41598-021-89249-6>
9. Jones, A., Haywood, J. M., Jones, A. C., Tilmes, S., Kravitz, B., & Robock, A. (2021). North atlantic oscillation response in geomip experiments g6solar and g6sulfur: why detailed modelling is needed for understanding regional implications of solar radiation management. *Atmospheric Chemistry and Physics*, 21(2), 1287–1304. <https://doi.org/10.5194/acp-21-1287-2021>
10. Shepherd, T. G. (2019). Storyline approach to the construction of regional climate change information. *Proceedings of the Royal Society*, 475(2225). <https://doi.org/10.1098/rspa.2019.0013>
11. Nordhaus, W. (2018). Evolution of modeling of the economics of global warming: changes in the DICE model, 1992–2017. *Climatic Change*, 149(4), 623–640. <https://doi.org/10.1007/s10584-018-2218-y>
12. Nordhaus, W. (2019). Economics of the disintegration of the greenland ice sheet. *Proceedings of the National Academy of Sciences*, 116(25), 12261–12269. Retrieved December 20, 2022, from <https://doi.org/10.1073/pnas.1814990116>. <https://www.pnas.org/content/116/25/12261.full.pdf>
13. Carlino, A., Giuliani, M., Tavoni, M., & Castelletti, A. (2020). Multi-objective optimal control of a simple stochastic climate-economy model. *IFAC-PapersOnLine* 53(2), 16593–16598. <https://doi.org/10.1016/j.ifacol.2020.12.786>. 21th IFAC World Congress.
14. Pindyck, R. S. (2017). The Use and Misuse of Models for Climate Policy. *Review of Environmental Economics and Policy*. <https://doi.org/10.1093/reep/rew012>

15. Botta, N., Jansson, P., & Ionescu, C. (2017). Contributions to a computational theory of policy advice and avoidability. *Journal of Functional Programming*, 27, 1–52. <https://doi.org/10.1017/S0956796817000156>
16. Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Hoboken, New Jersey: John Wiley & Sons.
17. Bertsekas, P. D. (1995). *Dynamic Programming and Optimal Control*. Belmont, Mass: Athena Scientific.
18. Ionescu, C., Klein, R. J. T., Hinkel, J., Kavi Kumar, K. S., & Klein, R. (2009). Towards a formal framework of vulnerability to climate change. *Environmental Modelling and Assessment*, 14(1), 1–16.
19. Ionescu, C. (2009). Vulnerability modelling and monadic dynamical systems. PhD thesis, Freie Universität Berlin. Retrieved December 20, 2022, from <https://d-nb.info/1023491036/34>
20. Brady, E. (2017). *Type-driven development in Idris*. Manning Publications Co., Shelter Island, NY 11964.
21. Brady, E. (2013) *Programming in Idris: A tutorial*. Retrieved December 20, 2022, from <https://www.idris-lang.org/docs/idris-tutorial.screen.pdf>
22. Ionescu, C., Jansson, P., & Botta, N. (2018). *Type theory as a framework for modelling and programming*. Springer, Cham, Switzerland.
23. Webster, M. D. (2008). Incorporating path dependency into decision-analytic methods: An application to global climate-change policy. *Decision Analysis*, 5(2), 60–75.
24. Chockler, H., & Halpern, J. Y. (2004). Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research*, 22(1), 93–115.
25. Halpern, J. Y. (2006) Causality, responsibility, and blame: A structural-model approach. In *Proceedings of the 3rd International Conference on the Quantitative Evaluation of Systems. QEST '06* (pp. 3–8). IEEE Computer Society, USA. <https://doi.org/10.1109/QEST.2006.9>
26. Halpern, J. Y. (2014). *Cause, responsibility, and blame: A structural-model approach*.
27. Intergovernmental Panel on Climate Change (IPCC). (2018). IPCC, 2018: Summary for Policymakers. Intergovernmental Panel on Climate Change (IPCC).
28. Otto, I. M., Donges, J. F., Cremades, R., Bhowmik, A., Hewitt, R. J., Lucht, W., Rockström, J., Allerberger, F., McCaffrey, M., Doe, S. S. P., Lenferna, A., Morán, N., van Vuuren, D. P., & Schellnhuber, H. J. (2020). Social tipping dynamics for stabilizing earth's climate by 2050. *Proceedings of the National Academy of Sciences*, 117(5), 2354–2365. Retrieved December 20, 2022, from <https://doi.org/10.1073/pnas.1900577117>. <https://www.pnas.org/content/117/5/2354.full.pdf>
29. Rockström, J., Steffen, W., Noone, K., Persson, A., Chapin, F.S., Lambin, E., Lenton, T. M., Scheffer, M., Folke, C., Schellnhuber, H., Nykvist, B., De Wit, C. A., Hughes, T., van der Leeuw, S., Rodhe, H., Sörlin, S., Snyder, P. K., Costanza, R., Svedin, U., Falkenmark, M., Karlberg, L., Corell, R. W., Fabry, V. J., Hansen, J., Walker, B., Liverman, D., Richardson, K., Crutzen, P., & Foley, J. (2009). Planetary boundaries: Exploring the safe operating space for humanity. *Ecology and Society*, 14(2), 32.
30. Heitzig, J., Kittel, T., Donges, J. F., & Molkenthin, N. (2016). Topology of sustainable management of dynamical systems with desirable states: from defining planetary boundaries to safe operating spaces in the earth system. *Earth System Dynamics*, 7, 21–50.
31. Barfuss, W., Donges, J. F., Lade, S., & Kurths, J. (2018). When optimization for governing human environment tipping elements is neither sustainable nor safe. *Nature Communications*, 9, 2354. <https://doi.org/10.1038/s41467-018-04738-z>
32. Botta, N., Jansson, P., & Ionescu, C. (2018). The impact of uncertainty on optimal emission policies. *Earth System Dynamics*, 9(2), 525–542. <https://doi.org/10.5194/esd-9-525-2018>
33. Braham, M., & van Hees, M. (2018) Voids or fragmentation: moral responsibility for collective outcomes. *The Economic Journal*, 128(612), 95–113. Retrieved December 20, 2022, from <https://doi.org/10.1111/eoj.12507>. <https://academic.oup.com/ej/article-pdf/128/612/F95/27015592/ejof95.pdf>
34. Heitzig, J., & Hiller, S. (2020). Degrees of individual and group-wise backward and forward responsibility in extensive-form games with ambiguity, and their application to social choice problems. Retrieved December 20, 2022, from <https://arxiv.org/abs/2007.07352>
35. Brede, N., & Botta, N. (2021). On the correctness of monadic backward induction. *Journal of Functional Programming*, 31, 26. <https://doi.org/10.1017/S0956796821000228>
36. Wadler, P. (1992). Monads for functional programming. In *Program Design Calculi, Proceedings of the NATO Advanced Study Institute on Program Design Calculi, Marktobendorf, Germany, July 28 - August 9, 1992* (pp. 233–264). https://doi.org/10.1007/978-3-662-02880-3_8
37. Botta, N., Brede, N., Jansson, P., & Richter, T. (2021). Extensional equality preservation and verified generic programming. *Journal of Functional Programming*, 31, 24. <https://doi.org/10.1017/S0956796821000204>
38. Botta, N., Jansson, P., Ionescu, C., Christiansen, D. R., & Brady, E. (2017). Sequential decision problems, dependent types and generic solutions. *Logical Methods in Computer Science*, 13(1). [https://doi.org/10.23638/LMCS-13\(1:7\)2017](https://doi.org/10.23638/LMCS-13(1:7)2017)
39. Bellman, R. (1957). *Dynamic programming*. Princeton University Press, Princeton, New Jersey.
40. Mitchell, T. M. (1997). *Machine Learning*, 1st edn. McGraw-Hill, Inc., USA.
41. Botta, N., et al. (2019–2021). *Papers*. Retrieved December 20, 2022, from <https://gitlab.pik-potsdam.de/botta/papers>
42. Sharpe, S., Mercure, J.-F., Vinales, J., Ives, M., Grubb, M., Pollitt, H., Knobloch, F., Nijssen, F. J. M. M. (2021). Deciding how to decide: Risk-opportunity analysis as a generalisation of cost-benefit analysis. Technical report, UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2021/03).
43. Carnap, R. (1995). *Introduction to philosophy of science*. Dover, New York.
44. Botta, N., et al. (2019–2021). *IdrisLibs2*. Retrieved December 20, 2022, from <https://gitlab.pik-potsdam.de/botta/IdrisLibs2>
45. Ross, D. (2021). Game theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2021 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2021/entries/game-theory/>
46. Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* (vol. 4). Selten254 International Journal of Game Theory 1975, 25–55
47. Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.
48. Hinze, R., & Löh, A. (2015). *Guide to lhs2TEX*. 18th International Workshop on Types for Proofs and Programs (TYPES 2011). Retrieved December 20, 2022, from <https://hackage.haskell.org/package/lhs2tex-1.19/src/doc/Guide2.pdf>
49. Ionescu, C., & Jansson, P. (2013). Testing versus proving in climate impact research. In *Proc. TYPES 2011. Leibniz International Proceedings in Informatics (LIPIcs)* (vol. 19, pp. 41–54). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany. <https://doi.org/10.4230/LIPIcs.TYPES.2011.41>
50. Howard, W. A. (1969). The formulae-as-types notion of constructions. In *To HB Curry: Essays on Combinatory Logic, Lambda*

Calculus and Formalism. Academic Press, London. Published only in 1980.

51. Wadler, P. (2015). Propositions as types. *Communications of the ACM*, 58(12), 75–84. <https://doi.org/10.1145/2699407>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.