Functional multi-objective optimization under uncertainty

Patrik Jansson (main PI)

1 Purpose and aims

The latest IPCC report clearly shows that avoiding unmanageable consequences of climate change requires a profound transformation of the ways we produce, transport and consume energy, commodities and services. But the report provides little clues on which concrete climate policies can promote and sustain such a transformation through the next decades. This is not very surprising: designing fair carbon taxes, and effective schemes for trading emission rights is difficult. And bans on cheap, emissions intensive technologies are difficult to agree upon, and if anticipated can lead to even higher global emissions. In a nutshell: a transformation towards a fossil fuel free society must itself be sustainable. It must reach the national climate goals but also ensure safety and prosperity of future generations. This requires an open, iterative decision process that integrates climate science with empirical data, domain-expert knowledge, etc. into stylized decision problems. These problems need to handle uncertain knowledge. But they also have to account for uncertainties in the decision process itself: decisions are often delayed or postponed. Doing so requires a holistic approach and a method that can deal with different kinds of uncertainty rigorously.

We aim to provide decision makers and advisors with tools for computer-aided decision making. These tools will allow them to

1. specify the most relevant uncertainties that affect a decision problem;
2. put forward different, possibly incompatible, assumptions for a decision problem;
3. deduce the (most likely, worst, or best) possible consequences of such assumptions on the climate system, on the economic system, and on societies;
4. discover and explore policies that fulfil specific requirements, for example, that are "best" for avoiding the worst.

Perhaps more crucially, we would like these methods to help climate policy advisors to explicitly put forward all the assumptions upon which a given policy rests and that in modelling approaches are often "buried" in the model. Thus, the focus of the application is on building tools for advisors and decision makers that are understandable, transparent and correct.

We will do this through specification, prototype implementation, and verification of multi-objective control problems under uncertainty, and by building domain-specific languages for exploring their solutions.

2 State-of-the-art

While the last IPCC AR6 Synthesis Report [1] leaves virtually no doubt that "*human activities . . . have unequivocally caused global warming*", mainly through "*unsustainable energy use, land use and land-use change, lifestyles and patterns of consumption and production*", it has become increasingly difficult for humanity to agree upon **policies** that avoid unmanageable consequences of climate change, not to mention to implement such policies.

There are two reasons: first, global players are increasingly in competition (or in open conflict) and the (real or perceived) exposure to the risks of climate change is different for different players. As a consequence, they find it increasingly difficult to agree on even how to value the risks of climate change, not to mention on how to share the burden of **avoiding** or **mitigating** such risks. The second reason why it is so difficult to negotiate and implement global climate policies is **uncertainty**: rationalizing decision making under deep uncertainty is difficult, both methodologically and computationally.

It is important to model the consequences of climate change on economies and societies under different greenhouse gas (GHG) emission scenarios and, conversely, the impact of societal transformations (e.g., of "green" transitions to decarbonized economies) on GHG emissions, but modelling is not enough: models are uncertain, and decisions are politically fraught and imperfectly implemented. It is plain that there is no "technical" solution to these difficulties (see "The Tragedy of the Commons" [2]) and yet policy makers need to negotiate and implement accountable decisions now: waiting for more reliable knowledge (or technological innovation) to become available is no longer an option.

3 Significance and scientific novelty

What is needed, and what we will provide, is a computer-aided methodology to **combine** the results obtained with these models (but also through expert elicitation and plain speculative assumptions, for example, about the likelihood that a country will implement a certain policy) in **consistent specifications** of simple but representative, **stylized** decision problems. (A well-known example of a stylized decision problem is the Prisoner's dilemma [3] which despite its simplicity can still help understanding climate problems involving "free-riding".)

Further, we need **trustable** computational methods for exploring specific regions of the solution space of such problems: what are policies that do not make any country worse off (and perhaps some of the countries better off)? What are decisions that ensure (at least, with a given confidence) that certain critical thresholds (tipping points) are not exceeded?

Mainstream integrated assessment models (IAMs) [4, 5] can be applied to answer these questions in a perfect world, one in which the response of the climate system to anthropogenic forcings is known, costs and benefits can always be quantified with certainty, decision makers are perfectly efficient and binding international agreements are always met.

Alas, a realistic approach towards decision making needs to account for the fact that our understanding of the climate response is far from perfect, that real costs do not always agree with the plans and that decisions are not always implemented with certainty, be this because of the intrinsic inertia of legislations or because some other global challenge (a pandemic or a war) has taken centre stage. We believe that, for **pragmatic** decision making, modelling needs to be integrated in a consistent theory and complemented with simpler and more realistic computational methods. What are meaningful requirements for a consistent theory of policy advice and for realistic methods for pragmatic decision making?

As a minimum, the theory should support dealing with different kinds of uncertainty in a rigorous, consistent manner. The focus on different kinds of uncertainty is crucial because, in climate policy advice, it is often difficult (unrealistic) to assign well defined probabilities to possible transitions. What is the probability that 80% of the Greenland ice sheets will melt over the next 1000 years, given a particular GHG emissions scenario? Nobody knows for certain and it is not even clear what the question means precisely. But perhaps it is possible to formulate this precisely and to come up with a realistic range of probabilities, rather than with a single figure.

Further, the theory should ideally be **verified**. In contrast to decision making in engineering and in the physical sciences, climate science and policies can hardly be verified empirically [6]. Not even global players like China or the USA can afford to perform large scale, carefully designed social experiments in order to assess the effectiveness of, say, carbon taxation schemes. And testing two or three carbon

taxation schemes on a couple of EU countries to find out which one would be best to adopt on a larger scale is simply impossible. And even if it were possible, it would take time that is not any longer at our disposal.

The only way to provide decision makers with some guarantees is through methods that, bar the validity of the assumptions that need to be put forwards to apply those method, have been verified to fulfil a precise specification. This is not a peculiarity of climate policy advice: advising governments on how to auction radio frequencies or internet domains faces similar problems [7], and related problems are also encountered in policy advice on matters of epidemics, financial markets, taxation, and security.

What about realistic methods for pragmatic decision making? We believe that policy advice that helps administrations and negotiators to **agree** upon climate policies needs to be underpinned by a verified theory but based on "fast and frugal" [8] methods.

We aim to provide decision makers and advisors with correct tools for computer-aided decision making. These tools will allow them to 1) specify the most relevant uncertainties that affect a decision problem; 2) put forward different, possibly incompatible, assumptions for a decision problem; 3) deduce the (most likely, worst, or best) possible consequences of such assumptions on the climate system, on the economic system, and on societies; 4) discover and explore policies that fulfil specific requirements.

We would like these methods to help climate policy advisors to explicitly put forward all the assumptions, for example, about discount factors or climate sensitivity, upon which a given policy rests and that in modelling approaches are often "buried" in the model [9]. In a nutshell, the focus of the application is on building tools for advisors and decision makers that are understandable, transparent and correct.

4 Preliminary and previous results

The PI has been involved in a long term collaboration with the Potsdam Institute for Climate Impact Research (PIK) on sustainability science and verified methods for optimal decision making in the context of climate impact research. This collaboration has lead to a sequence of papers which currently (March to June 2024) makes up the reading list for the PhD course "Functional Programming and Climate Impact Research" lead by the PI at Chalmers. The collaboration has also lead to a research programme (including three EU projects) which the current proposal aims to continue.

The starting point of this work was the PhD thesis by C. Ionescu [10] on vulnerability modelling with functional programming and later also with dependent types [11]. We then moved on to exploring and explaining the pros and cons of testing versus proving in climate impact research [12], expanding on earlier work on testing of polymorphic properties [13] using parametricity [14]. Next up we formalised concepts from economic modelling as examples of dependently-typed programming in scientific computing [15]. We worked in three different EU projects on developing a computational theory of Global Systems Science based on Domain-Specific Languages [16]. This lead to a first formulation of climate-relevant sequential decision problems using dependent types and generic programming [17] and extensions to a computational theory of policy advice and avoidability [18].

Directly on topic for the current proposal is our paper on the impact of three kinds of uncertainties on optimal emission policies [19]. Our most recent results in this direction explores responsibility under uncertainty: which climate decisions matter most [20]? Here we have applied the dependently typed, verified theory of policy advice [18] to understand how much decisions under uncertainty matter [20] in stylized climate decision problems. We proposed (and implemented a prototype for) a method to measure the relative importance of different decision steps. Extending this approach to more realistic problems requires trustable computational methods for multi-objective optimal control problems under uncertainty, the topic of the current proposal.

In 2022–2023, together with the Physics department at Chalmers, N. Botta and the PI have applied Bayesian optimization [21, 22], adaptive mesh refinement [23] and random walk methods to explore

Pareto optimal solutions for low-dimensional multi-objective optimal control problems. These techniques will be used, and extended, in the proposed project.

5 Project description

As mentioned above, the PI and N. Botta have contributed towards building a verified theory of policy advice and avoidability [18]. The theory was applied in [19] to understand the impact of uncertainties on optimal emission policies and in [20] to quantify responsibility in climate decisions under uncertainty. In this project, we want to build on this basis and extend both the theory and its range of applications. While the theory is general enough to cope with different kinds of **monadic** uncertainty, its focus is on policies that optimize a specific sum of the rewards associated with the possible trajectories of a (deterministic, non-deterministic, stochastic, etc.) system.

In climate policy advice, one is often interested in balancing the tradeoffs of policies for different, typically conflicting, goals. For example, one would like to avoid unmanageable consequences of climate change but also growing social inequality and perhaps other undesirable consequences of fast, irreversible societal transformations. In other words, climate policy advice is a **multi-objective control problem**. There is a system, a space of control policies, and objectives to be met. No control policy is best, but some are **Pareto-optimal** (see fig. 1) in that one can not improve one objective without worsening another [24].

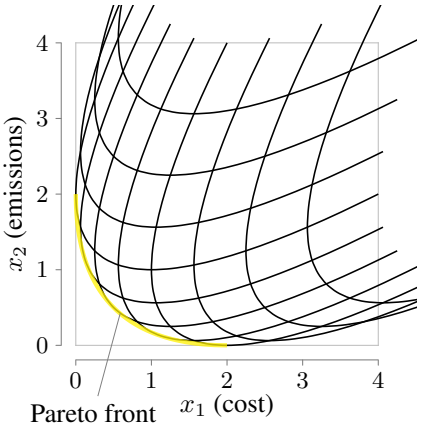


Figure 1: Let the objective function be $f : Y \to X$ with $Y = X = \mathbb{R}^2$. Each curve keeps one control (y_1 or y_2) fixed and varies the other continuously. The Pareto front is the set of Pareto-optimal points in the objective space X.

It is well known [25] that Pareto optimality is a weak notion which needs to be complemented with additional (stability, robustness, etc.) criteria and with methods for exploring and measuring different regions of the policy space with respect to such criteria [26].

The project aims to extend the policy advice and avoidability theory [18] and to develop and apply a toolkit of open-source computational methods for specifying multi-objective decision problems under monadic uncertainty and for exploring safe and stable solutions in close collaboration with domain experts from sustainability science. The project is designed around three main research questions and corresponding work packages:

- WP1: How to encode stylized decision processes?
- WP2: How to encode preferences and goals?
- WP3: How to explore alternative policies?

We do this through specification, prototype implementation, and verification of multi-objective control problems under uncertainty, and by building domain-specific languages for exploring their solutions.

5.1 Theory and method

WP1: How to encode stylized decision processes?

Decision problems in climate impact research are often finite-horizon sequential decision problems under uncertainty [27]. In the economic literature, these problems are often referred to as "intertemporal optimization" or "optimal planning" problems. When uncertainties are accounted for, this is typically done in probabilistic terms and measuring uncertainties with the expected value. The decision **process** underlying a problem is then specified in terms of a transition probability table or matrix. This is hardly compositional and yields low-level specifications which are difficult to understand, explain and modify.

A goal of WP1 is to support the specification of decision processes in diagrammatic form, through composable Bayesian Belief Networks (BBN) [28]. Bayesian belief networks are Directed Acyclic Graphs (DAGs) with finite conditional probability tables in the nodes, similar to the "causal" networks discussed in [29]. Augmented with annotations for specifying physical dimensions, units and informal semantics of the quantities involved, these probability tables are ideal candidates for specifying decision processes compositionally in a simple, consistent, and yet flexible way.

We have preliminary results (published as the "[...] which climate decisions matter most?" paper [20]) and in WP1 we will develop a small domain-specific language to help experts specify complex transition processes by "filling-in" and composing simple BBNs. The resulting DAGs are then functionalized and the resulting transition function becomes the dynamical system (or joint probability table) of the whole decision process. This guarantees that uncertainties are propagated consistently with the uncertainty model (through the monadic operators).

The second part of WP1 is the development of interfaces for specifying the most basic ingredients of climate decision problems: state and control spaces, reachability and viability predicates, and, perhaps most importantly, the relationship between states and available controls.

In most practical cases, the controls available at a given decision step depend on the state at that step. For example, in GHG emission processes, geo-engineering options for mitigating the effects of high GHG concentrations in the atmosphere might be available (e.g., because of binding agreements) only if the global mean temperature exceeds certain "tolerance" thresholds. And it is clear that the availability of expensive carbon extraction options might be conditional to budget and perhaps other constraints.

Close collaboration with domain experts [20, 30] has shown that obtaining a clear specification of the decision processes underlying a decision problem is a difficult and sometimes painful task that requires both scientists and advisors to think "out of the box". On the other hand, a clean process specification is a pre-condition for the precise study of any sequential decision process. The task of WP1 is to make the specification of decision processes as easy as possible but also to guarantee that such specifications are complete and consistent.

WP2: How to encode preferences and goals?

Decision processes become decision problems through the specification of the goals of decision making. This requires specifying the value of taking a sequence of decisions according to a given sequence of policy functions and starting from an initial state or, in the case of imperfect information, from a set of possible initial states. There are different ways of formulating such "value functions". A very common one is by measuring the value of each single "possible" trajectory that starts with the initial state (or with any of the possible initial states) through one or more **objective** functions [31].

The values associated with the possible trajectories then have to be aggregated for each objective: for example, avoid the worst possible consequences of climate change, ensure the availability of affordable energy, or avoid economic downturns. In stochastic decision problems, this aggregation is typically done via the expected value measure. This has a number of advantages, among others, that of being compatible with a well understood approach for computing optimal policy sequences (using dynamic programming

and backward induction [31–34]). But when dealing with outcomes which are possible albeit very unlikely and whose consequences are unacceptable, other measures are often more appropriate.

Sometimes, the different objectives can themselves be aggregated. Again, there is no canonical way of doing so. When the relative importance (or value) of the objectives is known and the value of possible trajectories is measured with numbers, fixing $n-1$ **prices** is a way of aggregating n objective functions. But sometimes measures are Boolean functions and aggregation is more naturally defined in terms of logical operations: conjunctions, disjunctions, etc.

The task of WP2 is to develop a small domain specific language for specifying the preferences and the goals of decision making in a way that is consistent with the uncertainties that affect the underlying decision problem. Again, we have found out that this is not a trivial task and that domain experts and practitioners are often not aware of the need (and of the possibilities) of putting forward precise specifications, but also of the consequences of inconsistent specifications of the goals of decision making on policies ("garbage in, garbage out").

WP3: How to explore alternative policies?

The aim of WP1 and WP2 is to allow combining existing knowledge (from model simulations, domain experts, empirical data) consistently into precise specifications of simple but representative multi-objective decision problems under uncertainty. The aim of WP3 is to exploit these specifications to give policy advisors, negotiators and practitioners trustable tools for exploring alternative solutions to these problems.

The emphasis here is on "alternative". As discussed in the State-of-the-art section, there is no "technical" solution to the problem of preventing human activities to drive the climate through and beyond critical thresholds that eventually will threaten the human civilization. In this context, the notions of optimality that are naturally associated with problem specifications are to be taken with a grain of salt. Optimal policies are typically non-unique and two optimal policies can be essentially different. The problem specifications themselves are necessarily freighted with assumptions and value judgments that often are, to say the least, questionable. And uncertainties and the horizon of decision making may have unexpected impacts on what are finally "best" climate policies.

In a nutshell: we want to provide decision makers with methods that do not hide the complexities and the ambiguities of decision making under uncertainty, but rather enable them to understand, explore and learn the logical and practical implications of these complexities. This requires, on the one hand, methods for computing "best" decisions for stylized decision problems and, on the other hand, methods for assessing the consequences of best decisions and comparing such consequences to those of other decision rules or policies.

A crucial requirement is that these methods are **trustable**. There are basically two methodologies for developing computational methods that can be trusted: testing and proving [12]. In WP3, we will crucially need both. We will leverage on our work done in [18, 34, 35] but also on specification-based testing with industrial standards [36].

Based on the stylized examples from WP1 and the DSL from WP2 we will provide a few case-studies which highlights the strengths and weaknesses of competing "best" (Pareto-optimal) decisions.

5.2 Time plan and implementation

WP1 (encode stylized decision processes) and WP2 (encode preferences and goals) will be active during the first phase of the project and until the end of the second project year. A risk here is the dependence on domain experts from sustainability science, but our long-term collaboration with PIK and the growing "Programming for the Planet" network (see section 7) will help us managing this risk.

During the first two project years, the focus of WP3 (explore alternative policies) will be on methods for verified backwards-induction, for exploring safe, Pareto-optimal policies under monadic uncertainty and for the statistical analysis of the possible consequences of policy sequences. This research will require assessing and comparing the effectiveness of Bayesian optimization, adaptive mesh refinement and random walk methods for exploring the control policy space and the development of methods for clustering possible trajectories.

In the third and fourth project year we will apply the methods to systematically specify and study full-fledged stylized decision problems. These will be selected from a number of smaller, partial problems that we will specify and study during the first two project years to exercise the methods developed in all work packages. A risk is that the decision problems grow too big to effectively handle with our methods, but our focus on stylized problems, where unnecessary detail is factored out, will help mitigate this risk.

5.3 Project organisation

The core team consists of the principal investigator (PI, Patrik Jansson, Chalmers) and senior scientist (Nicola Botta) from the Potsdam Institute for Climate Impact Research (PIK). Both will work at least 20% on this project throughout the four years. This will ensure supervision capacity and time to lead and contribute to the main tasks of the project.

We will recruit one PhD student working 80% on the project (and the remaining 20% on teaching, funded by other sources) to work on language design, modelling, implementation, documentation, and verification, supervised by the core team. We will also involve two student research assistants (SREs) at 20% each in parallel with their studies (funded by other sources). We have good experience with SREs (the PI has supervised nine over the years) and they help out mainly on the implementation and documentation side.

We will to report our main results in scientific journals and at conferences, but we also aim to post regular updates on the project homepage and share the code in public source code repositories.

6 Need for research infrastructure

We do not foresee any need for national or international research infrastructure other than already existing source code repositories.

7 International and national collaboration

The PI has been involved in a long term collaboration with the Potsdam Institute for Climate Impact Research (PIK, Germany) on sustainability science and verified methods for optimal decision making in the context of climate impact research. This connection will enable close collaboration with domain experts from sustainability science.

We have been in the PC of, and speaking at, the first edition of "Programming for the Planet (PROPL)[1]" as a SIGPLAN workshop affiliated with POPL'24. We are collaborating with Prof. A Madhavapeddy, Cambridge who has given a guest lecture at the "Functional Programming and Climate Impact Research (FPClimate)[2]" PhD course currently running at Chalmers (Mar–Jun 2024). The course helps spreading the word about the proposed research, recruiting new talent, and setting up new collaborations.

We have strong links to Prof. C Ionescu, Deggendorf Institute of Technology, Germany, a former PostDoc of the PI and co-author of several of the main publications leading up to this proposal. Prof. Ionescu is also part of the organizing team for the FPClimate course.

[1] The PROPL workshop homepage is at popl24.sigplan.org/home/propl-2024

[2] The FPClimate course homepage and code repository is at github.com/DSLsofMath/FPClimate.

The PI is a long term member, and vice-chair, of the IFIP working group 2.1 on Algorithmic Languages and Calculi. The group organizes regular work meetings to exchange ideas about the latest research in the programming languages area. Important collaborators from this group include C Ionescu, W Swierstra, and J Gibbons. During autumn 2024 the PI will be on a sabbatical visiting Prof. J Gibbons, Oxford and this will include at least two seminars related to the topic of this application (one in Oxford, one in Cambridge).

We also have ongoing local collaboration with the plasma physics group lead by Prof. T Fülöp at the Physics department, Chalmers on multi-objective optimization under uncertainty in the energy area. This provides an interesting future application area for the methods and tools developed in the currently proposed project.

Last, but not least, we benefit from collaboration inside the internationally strong group on functional programming and dependent types at Chalmers and University of Gothenburg. The very good reputation of the group also helps in attracting very strong candidates for PhD positions and MSc projects.

8 Societal value

In discourse about climate policies, a major misconception is that the goal of decision making is to achieve agreed climate targets, for example, holding "*the increase in the global average temperature to well below 2°C above pre-industrial levels*" or, better, limiting "*the temperature increase to 1.5°C above pre-industrial levels*" as stated in the 2015 Paris agreement. The truth is that climate policies cannot be informed by climate targets alone: indeed, it would be very "easy" (but very painful) to limit the temperature increase to below 1.5°C if humanity accepted draconian measures, like a massive reduction of the global population or of economic activities.

As convincingly pointed out in [29], a sound perspective towards climate decision making necessarily "*involves an acknowledgment that climate-relevant decisions, especially at the local scale, are not usually made on the basis of climate change alone but involve many other changing factors, most of which are highly uncertain*".

This project is a first step towards formulating and accounting for (what is known and what is unknown about) those other factors in a flexible and rigorous way. We recognise that climate science has provided (and will continue providing) crucial information for rationalising decision making in matters of climate policy. But rationalising such decision making requires "*moving the climate risk problem out of the domain of pure climate science*" [29].

In other words: the climate policy of any welfare nation should be a coherent, stable and genuine component of its governance that informs and permeates both its internal and its foreign policy. Computer science can contribute toward making climate policy advice more accountable, more consistent, more pragmatic and finally more understandable.

In a previous project, we have shown that notions of responsibility under uncertainty that are built on top of verified methods for decision making can contribute to clarify and answer crucial questions in climate policy analysis. For example, the question of whether it is worth fighting for early decarbonization measures when the likelihood that such decisions will actually turn into actions becomes very small. In the "Matter Most" paper [20] we have found, perhaps surprisingly, 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.

In the same project, we have also found that, while a strong concern (think of the "Fridays for future" movement) that decisions that are taken (or delayed) now will have severe consequences on the options that will (not) be available to upcoming generations is fully justified, the claim that current decisions matter more than decisions that will be taken by future generations stands on a less solid ground.

With the research programme outlined in this application, we aim at producing similar insights for decision problems in which the tradeoffs of decision making cannot be measured in terms of a single objective function and the outcomes of decisions are deeply uncertain. We clearly do not know what is the societal value of such insights. We hope that a better understanding of the consequences of multiple goals, uncertainties and imperfect knowledge in climate decision processes will foster a more rational and collaborative approach, for example in climate negotiations and international agreements but also in the civil societies.

But we are well aware that a war or a pandemic can radically change the perception of what is the societal value of a research programme. The one outlined in this application is the result of a long standing collaboration between the PI and PIK. In [29], the authors suggest that climate policy advice should rest on four core values: humility, integrity, transparency, and collaboration. Our research programme is informed by these values.

When it comes to the UN Sustainable Development Goals, this project is expected to mainly contribute to goal 13 (climate action) but also to goals 11 (sustainable cities) and 4 (quality education). By developing a toolkit for specifying multi-objective decision problems under uncertainty we will enable experts and policy-makers to make more informed decisions about the choice of climate actions and the development of sustainable cities. And by enabling the exploration of safe and stable solutions to sequential decision problems we help improve the quality of education, both of university students and of policy-makers (and perhaps even the general public).

References

[1] Intergovernmental Panel on Climate Change (IPCC). *AR6 Synthesis Report Climate Change 2023: Summary for Policymakers*. en. Ed. by P. Arias et al. https://report.ipcc.ch/ar6syr/pdf/IPCC_AR6_SYR_SPM.pdf. 2023.

[2] G. Hardin. "The Tragedy of the Commons". In: *Science* 162.3859 (1968).

[3] R. Axelrod. "Effective Choice in the Prisoner's Dilemma". In: *Journal of Conflict Resolution* 24.1 (1980). DOI: 10.1177/002200278002400101.

[4] W. Nordhaus. "Evolution of modeling of the economics of global warming: changes in the DICE model, 1992–2017". In: *Climatic Change* 149.4 (2018). DOI: 10.1007/s10584-018-2218-y.

[5] W. Nordhaus. "Economics of the disintegration of the Greenland ice sheet". In: *Proceedings of the National Academy of Sciences* 116.25 (2019). DOI: 10.1073/pnas.1814990116.

[6] V. Lucarini. "Towards a definition of climate science". In: *ArXiv Physics e-prints* (Aug. 2004).

[7] M. Caminati et al. "Sound auction specification and implementation". In: *EC '15 Proceedings of the 16th ACM Conference on Economics and Computation*. Ed. by T. Roughgarden, M. Feldman, and M. Schwarz. Association for Computing Machinery, June 2015. DOI: 10.1145/2764468.2764511.

[8] G. Gigerenzer, J. Czerlinski, and L. Martignon. "How Good are Fast and Frugal Heuristics?" In: *Decision Science and Technology: Reflections on the Contributions of Ward Edwards*. Ed. by J. Shanteau et al. Boston, MA: Springer US, 1999. DOI: 10.1007/978-1-4615-5089-1_6.

[9] R. S. Pindyck. *The Use and Misuse of Models for Climate Policy*. Review of Environmental Economics and Policy. 2017. DOI: 10.1093/reep/rew012.

[10] C. Ionescu. "Vulnerability Modelling and Monadic Dynamical Systems". PhD thesis. Freie Universität Berlin, 2009.

[11] C. Ionescu. "Vulnerability modelling with functional programming and dependent types". In: *Mathematical Structures in Computer Science* 26.01 (2016). DOI: 10.1017/S0960129514000139.

[12] C. Ionescu and P. Jansson. "Testing versus proving in climate impact research". In: *Proc. TYPES 2011*. Vol. 19. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl: LZI, 2013. DOI: 10.4230/LIPIcs.TYPES.2011.41.

[13] J. Bernardy, P. Jansson, and K. Claessen. "Testing Polymorphic Properties". In: *ESOP 2010*. Ed. by A. D. Gordon. Vol. 6012. Lecture Notes in Computer Science. Springer, 2010. DOI: 10.1007/978-3-642-11957-6_8.

[14] J. Bernardy, P. Jansson, and R. Paterson. "Proofs for free - Parametricity for dependent types". In: *J. Funct. Program.* 22.2 (2012). DOI: 10.1017/S0956796812000056.

[15] C. Ionescu and P. Jansson. "Dependently-typed programming in scientific computing: Examples from economic modelling". In: *24th Symposium on Implementation and Application of Functional Languages (IFL 2012)*. Ed. by R. Hinze. Vol. 8241. LNCS. Springer, 2013. DOI: 10.1007/978-3-642-41582-1_9.

[16] P. Jansson, C. Ionescu, and N. Botta. "Towards a Computational Theory of GSS: a Case for Domain-Specific Languages". In: *Third Open Global Systems Science Conference, http://global-systems-science.org/3-open-gss-conference/*. 2014.

[17] N. Botta et al. "Sequential decision problems, dependent types and generic solutions". In: *Logical Methods in Computer Science* 13.1 (2017). DOI: 10.23638/LMCS-13(1:7)2017.

[18] N. Botta, P. Jansson, and C. Ionescu. "Contributions to a computational theory of policy advice and avoidability". In: *J. Funct. Program.* 27 (2017). DOI: 10.1017/S0956796817000156.

[19] N. Botta, P. Jansson, and C. Ionescu. "The impact of uncertainty on optimal emission policies". In: *Earth System Dynamics* 9.2 (2018). DOI: 10.5194/esd-9-525-2018.

[20] N. Botta et al. "Responsibility Under Uncertainty: Which Climate Decisions Matter Most?" In: *Environmental Modeling & Assessment* (2023). DOI: 10.1007/s10666-022-09867-w.

[21] B. Shahriari et al. "Taking the Human Out of the Loop: A Review of Bayesian Optimization". In: *Proceedings of the IEEE* 104.1 (2016). DOI: 10.1109/JPROC.2015.2494218.

[22] I. Pusztai et al. "Bayesian optimization of massive material injection for disruption mitigation in tokamaks". In: *Journal of Plasma Physics* 89.2 (2023). DOI: 10.1017/S0022377823000193.

[23] W. Zhang et al. "AMReX: Block-structured adaptive mesh refinement for multiphysics applications". In: *Journal of HPC Applications* 35.6 (2021). DOI: 10.1177/10943420211022811.

[24] A. Carlino et al. "Multi-objective optimal control of a simple stochastic climate-economy model". In: *IFAC-PapersOnLine* 53.2 (2020). DOI: 10.1016/j.ifacol.2020.12.786.

[25] D. Ross. "Game Theory". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Fall 2021. Metaphysics Research Lab, Stanford University, 2021.

[26] P. Żebrowski et al. "Sharing the Burdens of Climate Mitigation and Adaptation: Incorporating Fairness Perspectives into Policy Optimization Models". In: *Sustainability* 14.7 (2022). DOI: 10.3390/su14073737.

[27] M. D. Webster. "Incorporating Path Dependency into Decision-Analytic Methods: An Application to Global Climate-Change Policy". In: *Decision Analysis* 5.2 (2008).

[28] T. M. Mitchell. *Machine Learning*. 1st ed. USA: McGraw-Hill, Inc., 1997. ISBN: 0070428077.

[29] T. G. Shepherd. "Storyline approach to the construction of regional climate change information". In: *Proc. R. Soc.* 475.2225 (2019). DOI: 10.1098/rspa.2019.0013.

[30] M. Martínez Montero et al. "SURFER v2.0: a flexible and simple model linking anthropogenic CO_2 emissions and solar radiation modification to ocean acidification and sea level rise". In: *Geoscientific Model Development* 15.21 (2022). DOI: 10.5194/gmd-15-8059-2022.

[31] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Hoboken, New Jersey: John Wiley & Sons, 2014.

[32] R. Bellman. *Dynamic Programming*. Princeton, New Jersey: Princeton University Press, 1957.

[33] P. Bertsekas D. *Dynamic Programming and Optimal Control*. Belmont, Mass.: Athena Scientific, 1995.

[34] N. Brede and N. Botta. "On the correctness of monadic backward induction". In: *Journal of Functional Programming* 31 (2021). DOI: 10.1017/S0956796821000228.

[35] C. Ionescu, P. Jansson, and N. Botta. "Type Theory as a Framework for Modelling and Programming". In: *ISoLA 2018 LNCS Proceedings*. Springer, 2018. DOI: 10.1007/978-3-030-03418-4_8.

[36] J. Hughes. "Software Testing with QuickCheck". In: *CEFP 2009, Revised Selected Lectures*. Springer, 2010. DOI: 10.1007/978-3-642-17685-2_6.