THE BULLETIN OF SYMBOLIC LOGIC
Volume 29, Number 2, June 2023

REVIEWS

The Association for Symbolic Logic publishes analytical reviews of selected books and articles in the field of symbolic logic. The reviews were published in ***The Journal of Symbolic Logic*** from the founding of the JOURNAL in 1936 until the end of 1999. The Association moved the reviews to this BULLETIN, beginning in 2000.

The Reviews Section is edited by Graham Leach-Krouse (Managing Editor), Albert Atserias, Mark van Atten, Clinton Conley, Johanna Franklin, Dugald Macpherson, Antonio Montalbán, Valeria de Paiva, Christian Retoré, Marion Scheepers, and Nam Trang. Authors and publishers are requested to send, for review, copies of books to *ASL, Department of Mathematics, University of Connecticut, 341 Mansfield Road, U-1009, Storrs, CT 06269-1009, USA.*

KATIE STEELE AND H. ORRI STEFÁNSSON. ***Beyond Uncertainty: Reasoning with Unknown Possibilities***. Elements in Decision Theory and Philosophy. Cambridge University Press, Cambridge, UK, 2021, 110 pp.

Consider the predicament of a policy-maker who must decide whether to implement Solar Radiation Management (SRM) on a global scale even though she is unaware of some of the potential consequences. Ideally, we would like Bayesian decision theory to be able to help the policy-maker with her decision. However, the standard prescriptions apply only to logically omniscient, fully aware agents. The uncertainty introduced by the policy-maker's limited awareness goes beyond the uncertainty that standard Bayesian decision theory can handle. Or so accepted wisdom has it.

In ***Beyond Uncertainty***, Katie Steele and H. Orri Stefánsson challenge this accepted wisdom. The book begins with the presentation of the standard (Bayesian) model of rational preference and belief that Steele and Stefánsson expand upon throughout. As is typical of philosophers, they prefer a Jeffrey-style framework upon which propositions are the objects of both beliefs and desires, and rational agents' preferences over outcomes are constrained by the axioms of subjective expected utility (SEU) theory.

In Section 2 this framework is extended to accommodate sequential decision-making and, in Section 3, limited awareness (or, more specifically, anticipated *awareness growth*) is situated within this extended framework. Since it has standardly been assumed that agents are fully aware of all the possibilities that are relevant to their decisions, some additional formal machinery is required to achieve this. To this end, Steele and Stefánsson introduce the notion of an 'awareness context' which allows for the characterisation of a *subjective* possibility space (i.e., a set of possibilities that is dependent on an agent's state of awareness). More formally, they define an agent's awareness context as the set, $\mathbf{X}$, of basic propositions that the agent is aware of. Next, they define possibilities as truth functions, ω_i, over the basic propositions in $\mathbf{X}$, and an agent's real set of possibilities, $\mathbf{W_X}$, as the subset of the putative set of possibilities that contains only those possibilities that the agent deems consistent. From here, each basic proposition X_i is associated with the collection of $\omega_i \in \mathbf{W_X}$ for which X_i is true; these collections generate a Boolean algebra of sets, $\mathcal{F}_\mathbf{X}$, in the usual way.

1079-8986/23/2902-0006
DOI:10.1017/bsl.2023.11

Table 1. The policy-makers decision

	E_1	...	E_n	??
f	$f(E_1)$	...	$f(E_n)$	$f(??)$
g	$g(E_1)$	...	$g(E_n)$	$g(??)$

Against this backdrop, awareness growth is characterised as the enlargement of an agent's subjective possibility space. The idea is that, upon awareness growth, there is a shift in awareness context from $\mathbf{X}$ to $\mathbf{X}^+$ where $\mathbf{X}^+ = \mathbf{X} \cup \mathbf{X_j}$ and $\mathbf{X_j}$ is the set of all basic propositions $X_j \notin \mathbf{X}$ that the agent becomes aware of. This in turn corresponds to a shift from $\mathbf{W_X}$ to $\mathbf{W_{X^+}}$ and from $\mathcal{F}_\mathbf{X}$ to $\mathcal{F}_{\mathbf{X}^+}$. Intuitively, one might think of awareness growth as a learning event. However, it should be apparent that it is distinct from standard (Bayesian) learning events wherein an agent gains evidence for a familiar proposition and the space of possibilities shrinks; instead, the space of possibilities expands. Due to this difference, conditionalisation does not apply in cases of awareness growth. Thus, two questions are left open: how should a rational agent represent anticipated awareness growth in their decision problem? And how should they revise their credences upon awareness growth?

The formal machinery that Steele and Stefánsson introduce in Section 3 turns out to be important for the answer that they provide to the second question. In particular, it drives the argument presented across Sections 4 and 5 which says that there is no general (conservative) norm of rationality that governs credence revision upon awareness growth. This argument challenges the popular view in the literature which upholds the norm of Reverse Bayesianism (E. Karni and M.-L. Vierø, *'Reverse Bayesianism': A choice-based theory of growing awareness.* ***American Economic Review,*** vol. 103 (2013), no. 7, pp. 2790–810). Roughly, Reverse Bayesianism can be interpreted as saying (in Steele and Stefánsson's terminology) that, upon awareness growth, a rational agent should update their credences in such a way that they preserve the probability *ratios* of nonnull events in their original awareness context (i.e., if s_1 is twice as likely as s_2 in the original awareness context, s_1 must be twice as likely as s_2 in the new awareness context). This proposal is ultimately rejected for two reasons. First, Reverse Bayesianism seems to produce counterintuitive results when the propositions that an agent becomes aware of are *evidentially relevant* to the comparison of propositions of which the agent was already aware. Second, since, on Steele and Stefánsson's picture, the set of possibilities associated with any basic proposition X_i is constituted by truth functions over the basic propositions in $\mathbf{X}$ and awareness growth involves a revision of $\mathbf{X}$, it therefore also involves a revision of the set of possibilities associated with any basic proposition X_i. As such, on this picture *all* credences may be affected by awareness growth in some way and any general conservative norm of credence revision, such as Reverse Bayesianism, is misguided.

Whilst Section 5.4 contains a positive upshot to this largely negative argument, the heart of Steele and Stefánsson's positive proposal lies in their response to the first question, which is spelled out in Sections 6 and 7. Here they first propose replacing maximally specific states with 'events' which describe those contingencies that the agent is aware of in as much detail as possible but which may nevertheless fail to fully determine the outcomes of the available acts. Second, they propose incorporating a *subjective* catch-all which signifies "some abstract proposition standing in for a broad class of contingencies that the agent thinks she may later be in a position to concretize" (p. 74).

If we return to the example of the policy-maker who must decide whether to implement SRM, the idea is to specify the decision problem as follows (Table 1):

Where f and g are the available acts (perhaps 'implement SRM' and 'abstain'), E_1 to E_n are the most fine-grained mutually inconsistent events the policy-maker can think of, and '??' is a subjective catch-all standing in for those contingencies that she suspects she is unaware of.

Represented like this, Steele and Stefánsson argue that the policy-maker's predicament is somewhat unremarkable and (if certain basic conditions are met) we can treat her as an EU maximiser just like any ordinary reasoner. However, they go on to canvas two norms of rationality—'Awareness Reflection' and 'Preference Awareness Reflection'—which they think should constrain the synchronic credences and desires (respectively) of agents like the policy-maker who anticipate their awareness will grow in rather specific ways.

Whilst the positive proposal spelled out in Sections 6 and 7 leaves several questions open, Steele and Stefánsson successfully lay the foundations for others working within normative decision theory and related areas of economics and computer science to take up these questions and continue the work of characterising the reasoning of rational, but less-than-fully aware, agents.

MAGDALEN ELMITT
University of Sydney, Sydney, Australia. E-mail: magdalen.elmitt@gmail.com

JAN KRAJÍČEK, ***Proof Complexity***, Encyclopedia of Mathematics and Its Applications, no. 170, Cambridge University Press, Cambridge, UK, 2019, xvi + 516 pp.

The book presents the field of proof complexity in its full breadth and depth. It starts historically, tracing proof complexity to foundational questions of mathematical logic, and ends with a question about proof complexity's nature: "what are the intrinsic reasons that some formulas are hard to prove? Can the proof complexity of some formulas be traced to the computational complexity of associated computational tasks?" (p. 477)

The central goal of proof complexity is to prove lower bounds to the size of proofs in various (propositional) proof systems. To date no superpolynomial lower bounds are known for standard textbook systems, called *Frege*, given by finitely many inference rules. However, already lower bounds for weak proof systems are well-motivated from a computer science perspective for their application to algorithm analysis. Systems around Resolution are related to Sat solvers, algebraic systems like Nullstellensatz or Polynomial Calculus to ideal membership algorithms, and semi-algebraic systems like Sherali–Adams or Sum-of-Squares to linear or semidefinite programming. While the combinatorially inclined research in this direction forms the "rudiments from which proof complexity can grow" (p. 473), it uses somewhat ad hoc methods tackling specific tautologies and proof systems. The book aims to presents "proof complexity as a whole entity rather than as a collection of various topics held together loosely by few notions. The frame that supports it is logic." (p. 4)

The gem of proof complexity is a (weakly) exponential lower bound on the size of bounded-depth Frege proofs of tautologies expressing the pigeonhole principle. Being *bounded-depth* means that the proof operates with formulas of some fixed $\wedge/\vee$-alternation rank. This goes back to Ajtai's 1988 article "The complexity of the pigeonhole principle" and "opened completely new vistas, showing that proof complexity is part of a much larger picture and that it does not need to be just a finitary proof theory" (p. 184). Ajtai gave a forcing-type construction of an expansion of a cut of a nonstandard model of true arithmetic by a bijection between $n+1$ and n for some nonstandard n in such a way that induction for bounded formulas is preserved. This implies the proof lower bound due to the correspondence of bounded-depth Frege and arithmetics with bounded induction.

That a bounded arithmetic T *corresponds* to a proof system P means that (1) P has short proofs of propositional translations of universal consequences of T and (2) T proves the soundness of P. By (1), lower bounds on P-proofs imply independence from T, and this explains a central motivation from mathematical logic (p. 37): understanding independence

https://doi.org/10.1017/bsl.2023.11 Published online by Cambridge University Press