Gödel-Turing barriers

for microscopic descriptions of the macroscopic world

January 31, 2023

The last decade has seen a plethora of new *undecidability* results in physics, especially in the field of quantum theory. Viewed from a distance, these can be divided into two categories of decision problems: those that somehow become unpredictable over an unbounded time domain, and those in which spatially bounded building blocks, when assembled, lead to undecidable global properties. In particular, the latter class raises the question whether there could be a fundamental barrier for microscopic descriptions of the macroscopic world, which ultimately originates in the foundational works of Gödel [Göd31] and Turing [Tur36].

In this piece, we will revisit the abundant literature of recent undecidability results from a bird's eye view. Instead of going into the very different physical details, we will elaborate their common rationale: a reduction from the halting problem that can be made more specific by recourse to Gödel's 2$^{\text{nd}}$-incompleteness theorem and more practically relevant by appeal to complexity-theoretic arguments. We will discuss robustly emergent macroscopic effects on uncomputable scales and in the end have earned a delicious hot beverage.

A recipe for unprovable propositions

Mathematical undecidability results come in two flavors: those in which undecidable means *uncomputable* in the sense of Turing, and those in which it means *unprovable* in the spirit of Gödel. Almost all undecidability results that are related to physics are based on the uncomputability of Turing's halting problem. However, as we will see a few paragraphs further down, they too lead to unprovable propositions, with the help of Gödel's 2$^{\text{nd}}$ incompleteness theorem.

The direct way from a formal system, such as ZFC, to an unprovable proposition with a physical interpretation without a detour through Turing's garden appears to be difficult. The known set-theoretic *forcing* techniques add 'generic' sets that tend to be elusive from a physical perspective. In spite of this, a notable attempt on this direct path has been made in [dCDdB90, DD12].

In the following, we will take a bird's eye view of the way unprovable propositions with a physical interpretation are obtained from Turing's halting problem. The first and main step is a reduction from the halting problem. In principle, the halting problem could be replaced by any other undecidable decision problem, but in practice, it seems to be the only bearing root. Formally, we start with a subset $L \subset \mathbb{N}$ that is non-recursive, i.e. for which the decision problem "$x \in L$?" is uncomputable. In the lingo of theoretical computer science, L specifies a *language* and the standard choice for L is the subset of enumerated Turing machines that halt on the empty input string. The centerpiece of every 'Turing-style undecidability paper' is then a computable mapping that assigns to every $x \in \mathbb{N}$ a logical sentence S_x in a formal system $\mathcal{F}$ that is equipped with an (in our case ultimately physical) interpretation such that

$$x \in L \quad \Longleftrightarrow \quad S_x \text{ is true.} \tag{1}$$

That is, there is a physically interpretable proposition on the right of Eq.(1) that is related in a computable manner to an instance x of an uncomputable decision problem. In the results collected below and in the references, each S_x is a statement about an individual physical system that is specified by x, like "x has a spectral gap", "x thermalizes", "x reaches a state with a particular property", etc..

The mapping $x \mapsto S_x$ is, of course, where all the work lies and where the wealth of results branches out in different directions.

Before we look at this mapping in more detail, let us understand the implications of Eq.(1) for provability. To this end, we assume that the formal system $\mathcal{F}$ is *effectively axiomatized*. This means that syntactic correctness, membership in the set of axioms, and validity of a proof should be algorithmically decidable. As a consequence, we obtain that it is in principle possible to recursively enumerate all provable sentences of $\mathcal{F}$ by looping over all potential proofs of increasing length and eliminating invalid ones. We also assume that $\mathcal{F}$ is *sound*, i.e. that everything provable in $\mathcal{F}$ is true. Now suppose that for all (but finitely many) x either S_x or $\neg S_x$ would be provable within $\mathcal{F}$. Then, since we can recursively enumerate all theorems of $\mathcal{F}$, there would be an algorithm to decide "$x \in L$?" — in contradiction to our choice of L. Hence, with $\mathcal{S} := \{S_x\}_{x\in\mathbb{N}}$ we get:

$$\mathcal{S} \text{ contains infinitely many propositions unprovable in } \mathcal{F}. \tag{2}$$

The unspecific infinity in (2) is basically the price one has to pay for resorting to computability theory — after all, finite decision problems are always decidable. However, by invoking Gödel's 2^{nd} incompleteness theorem one can, in principle, pinpoint a *specific* unprovable proposition in $\mathcal{S}$ if L was indeed the set of halting Turing machines. For this purpose, consider a Turing machine that loops over all provable sentences of $\mathcal{F}$ and halts if it has found an inconsistency. Let $x = c$ be the number of such a Turing machine. Then $c \in L$ if and only if $\mathcal{F}$ fails to be consistent. Since the latter cannot be proven within a sound and thus consistent $\mathcal{F}$ due to Gödel's 2^{nd} incompleteness theorem (tacitly assuming that $\mathcal{F}$ is rich enough) we obtain that

$$S_c \text{ is not provable in } \mathcal{F}. \tag{3}$$

In [YA16] an explicit 7910-state Turing machine was constructed that halts if and only if ZFC is inconsistent. Plugging this into any specific mapping $c \mapsto S_c$ then would lead to an explicit sentence that is unprovable in ZFC. This argument was laid out in [Cub21] for the case where each S_x is a statement about a spectral gap of a quantum spin system, but it is equally valid for all undecidability results that originate from the halting problem. And there are plenty ...[1]

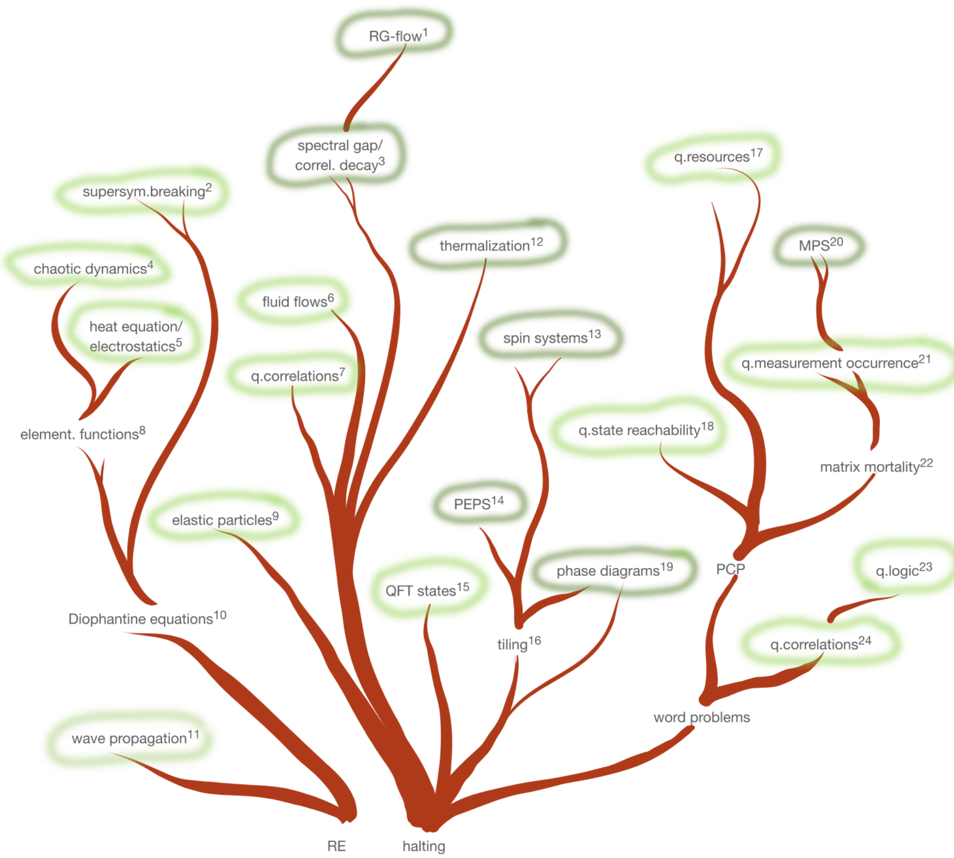

Figure 1: The tree of undecidable problems. Each branch connecting two classes of problems symbolizes a reduction argument used to prove the undecidability of the problem further up from the one below. Decision problems that have a deliberate physical interpretation are depicted in green. Darker green indicates problems asking for macroscopic properties of microscopically given systems. "q." stands for *quantum*. References are collected in the footnote.[1]

[1] Fig.1. displays the relation between the following references: [1] [WOC22], [2] [Tac22], [3] [CPGW22, CPGW15, BCLPG20], [4] [dCD91], [5] [CCSS94], [6] [Tao19, CMPSP21], [7] [JNV+20], [8] [Ric68, Wan74], [9] [FT82, Moo90, Moo91], [10] [Mat70, DPR61], [11] [PER81], [12] [SM21], [13] [Kan90], [14] [SMG+20], [15] [Kom64], [16] [Wan61, Ber66, Rob71], [17] [SS21], [18] [WCPG11], [19] [BCW21], [20] [KGE14, DlCCC+16], [21] [EMG12], [22] [CHHN14], [23] [Fri21], [24] [Slo20].

The undecidability-tree

The halting problem is the root of almost all provably 'undecidable' problems that have a physical interpretation. The reduction from the halting problem is, however, usually not done in one step (and if it is, that 'step' can easily cover dozens or a hundred of pages, as in [CPGW22] or [JNV$^+$20]). Typically, the map $L \to \mathcal{S}$ looks more like $L \to \mathcal{A} \to \mathcal{B} \to \cdots \to \mathcal{S}$ involving a chain of subsequent reductions. In this way, the set of existing results gets the structure of a *tree* with the physically interpretable undecidability results on the leaves or outer branches. The root is given by the halting problem or, more generally, by the set RE of *recursively enumerable* languages, of which the halting problem is the 'hardest' when viewed as decision problems. In the words of complexity theory, the halting problem is RE-*complete*.

Further up, the tree of undecidable problems has some prominent branches, three of which stand out as they are frequently used as intermediate reduction steps:

- *Diophantine equations*: these are polynomial equations in several integer-valued variables — like the equations that appear in Fermat's last theorem. A deep theorem by Matiyasevich [Mat70], Davis, Putnam and Robinson [DPR61] shows that RE can be identified with the class of *Diophantine sets*, which are projections of the solution sets of Diophantine equations.
- *Word problems*: theses exist for groups, semi-groups, (semi-)Thue-systems and other structures. Their common feature is the question whether two words formed from a given alphabet have the same 'meaning', i.e., fall into the same equivalence class. A particularly popular type of questions of this kind is *Post's correspondence problem* (PCP).
- *Tiling problems* ask whether it is possible to tile the plane with a certain set of tiles and rules for their arrangement.

The tree of undecidable problems is outlined in Fig.1. Needless to say, there are various forms of arbitrariness in this illustration, and as with any picture of a healthy tree, the number of visible leaves is not exhaustive.

The undecidability-tree has hard finite roots

How should we interpret the abundance of 'physically interpretable' undecidability results? Does anything of relevance remain if we remove some layers of idealization? A central and criticizable assumption is always a form of infinity that appears in each of the problem statements: infinite spatial extension, infinite time-scale, or infinite precision (or more than one of these). Climbing down the tree, we find the root of these infinities in the unbounded runtime of the Turing machine — an assumption that seems to be necessary to make the argument work. So what if we replace these infinities with something finite?

Then the considered problems may become decidable in principle but there is good reason to believe that many of them remain undecidable in practice:

The *time-bounded* version of the halting problem for (non-deterministic) Turing machines is one of the prime examples of the hardest problems within the notorious complexity class NP, i.e. it is NP-complete.[2] Thus, if the reductions along the branches are efficiently computable (which they typically are), then the physically interpretable problem further up the tree will inherit the difficulty of the root. In other words, while the finite analogues of the statements in $\mathcal{S}$ might become provable, finding their proofs will remain intractable (unless there is a major collapse in complexity theory). A more detailed elaboration of this line of thought can be found in [KvdES$^+$22].

We might go one step further, though, and shed some light on another potential shortcoming of this argument: after all, NP-completeness is a worst-case concept, which does not necessarily mean that instances of the considered problem are hard on average or typically hard. For example, the Hamiltonian path problem, while being NP-complete, takes only linear time on average [GS87]. In a similar vein, the knapsack problem is NP-complete but tractable in practice [Sha82].

However, the roots of the undecidability-tree, when cut down to a finite length, seem to be hard in a more robust way — they all appear to be *average-case* NP-*hard*. This has been proven for the bounded versions of the halting and tiling problem [Lev86], for word problems for groups [Wan95], Thue systems [WB95] and PCP [Gur91] and indicated for Diophantine problems in [VR92]. More recently, the bounded version of the halting problem was shown to be *strongly generically* NP-*complete* [MU16].

Beyond the analysis of worst-case complexity, reductions turn out to be considerably more subtle and difficult to prove. Hence, the above-mentioned average-case results do not automatically or easily spread further up the tree. The roots, however, seem to be particularly hard, even when cut to finite length and it would be strange if some of this did not also show up in the leaves.

Emergent effects on uncomputable scales

It is time to open the stage to the question of the extent to which the results discussed above pose an obstacle to the derivation of macroscopic phenomena from microscopic laws. In other words, what type of *emergence* do they indicate?

For this discussion, we will focus on those undecidability results that deal with such a spatial (rather than a temporal) problem. The *spectral gap problem* addressed in [CPGW22, CPGW15, BCLPG20, Cub21] for quantum spin lattices with local, translationally invariant interactions is one of these examples.

Undecidability results of this kind display high-level (macroscopic) phenomena that are only deducible from low-level laws if the scale of observation is

[2] In fact, it is even EXP-complete when the time-bound is encoded in binary rather than in 'unary'. However, this property is more difficult to preserve along reductions.

fixed (or bounded) in advance. However, if we ask: "is there a scale at which the phenomenon occurs?" the microscopic laws can be principally insufficient for an answer. More importantly, the distinction between the bounded and the open-ended case fades away if we invoke complexity-theoretic assumptions (such as $\mathsf{P} \neq \mathsf{NP}$ or its average-case counterpart $\mathsf{AvgP} \neq \mathsf{DistNP}$) under which 'intractable in practice' fades into 'intractable in principle'.

In this context, it is worth mentioning that the undecidability of the spectral gap aligns well with the notorious hardness of this problem faced by practicing mathematical physicists in even the simplest physical models. Take for example the antiferromagnetic spin-1 1D-Heisenberg model. In this case, the interaction between neighboring spins is described by a 9×9 matrix that has only 12 non-zero entries all of which are ± 1. The conjecture that a quantum system of spins interacting in this way on a spin-chain has a spectral gap is known as *Haldane conjecture* [Hal83, LWMA17]. Despite the blatant simplicity of the model, a proof of this conjecture is not in sight even after 40 years. The numerical evidence for this conjecture is currently limited to systems of a few hundred spins. But what if there are 10^5 or 10^{10}? The undecidability results show that extrapolations from smaller to larger scales can be unreliable, regardless of how large the 'small' scales were. That even relatively simple systems can abruptly change their properties at arbitrarily (and in general uncomputably) large scales, was exemplified in [BCL$^+$18].

The *size-driven phase transitions* of [BCL$^+$18] display a drastic change in the low energy properties of a 2D quantum spin model: below a lattice size N, the ground state is entirely classical and separated by a constant spectral gap. Above N, however, the low energy sector exhibits extremal quantum properties. The main point is that it is generally undecidable whether such an N exists, and that it can be astronomically large even for relatively simple systems ($N > 10^{36534}$ for spin dimensions ≤ 10 in the examples of [BCL$^+$18]).

In the discussion of *emergence* and *reductionism*, the question is often whether high-level laws can be deduced from low-level laws. The consequence of the discussed undecidability results is, however, that in these cases *there are no high-level laws.* Suppose there were: then we could add them to our formal system and run the exact same argument again — Turing's halting problem will still be undecidable and Gödel's 2$^{\text{nd}}$ incompleteness theorem will still be valid.[3] Furthermore, hinging on complexity-theoretic assumptions, this picture may not change if we bound the scale of observation: in some cases, the essentially only way of finding out whether a certain phenomenon will occur at a scale N of interest will be to go through the full computation or simulation up to that scale — no reliable heuristics or high-level law, no shortcut.

If you didn't arrive at the end of this piece via a shortcut, but made it here the hard way, you definitely deserve a delicious hot beverage. Go, get one, and enjoy ...

[3] This assumes, of course, that we stay on realistic grounds, away from non-recursive axiomatic systems and hypercomputation.

References

[BCL+18] Johannes Bausch, Toby S Cubitt, Angelo Lucia, David Perez-Garcia, and Michael M Wolf. Size-driven quantum phase transitions. *Proceedings of the National Academy of Sciences*, 115(1):19–23, 2018.

[BCLPG20] Johannes Bausch, Toby S. Cubitt, Angelo Lucia, and David Perez-Garcia. Undecidability of the spectral gap in one dimension. *Physical Review X*, 10(3):031038, 2020.

[BCW21] Johannes Bausch, Toby S. Cubitt, and James D. Watson. Uncomputability of phase diagrams. *Nature Communications*, 12(1):452, Jan 2021.

[Ber66] R. Berger. *The Undecidability of the Domino Problem.* Memoirs ; No 1/66. American Mathematical Society, 1966.

[CCSS94] Cristian Calude, Douglas I. Campbell, Karl Svozil, and Doru Stefanescu. Strong determinism versus computability. In *The Foundational Debate. Complexity and Constructivity in Mathematics and Physics. Vienna Circle Institute Yearbook, Vol. 3, ed. by Werner DePauli Schimanovich, Eckehart Koehler and Friedrich Stadler. Kluwer, Dordrecht, Boston, London, 1995, p.115-131*, pages 115–131, 1994, quant-ph/9412004.

[CHHN14] Julien Cassaigne, Vesa Halava, Tero Harju, and François Nicolas. Tighter undecidability bounds for matrix mortality, zero-in-the-corner problems, and more. *CoRR*, 2014, 1404.0644. arxiv:1404.0644.

[CMPSP21] Robert Cardona, Eva Miranda, Daniel Peralta-Salas, and Francisco Presas. Constructing turing complete euler flows in dimension 3. *Proceedings of the National Academy of Sciences*, 118(19):e2026818118, 2021.

[CPGW15] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. Undecidability of the spectral gap. *Nature*, 528:207, December 2015.

[CPGW22] Toby Cubitt, David Perez-Garcia, and Michael M. Wolf. Undecidability of the spectral gap. *Forum of Mathematics, Pi*, 10:e14, 2022.

[Cub21] Toby S. Cubitt. A note on the second spectral gap incompleteness theorem, 2021. arXiv:2105.09854.

[dCD91] N. C. A. da Costa and F. A. Doria. Undecidability and incompleteness in classical mechanics. *International Journal of Theoretical Physics*, 30(8):1041–1073, 1991.

[dCDdB90] N. C. A. da Costa, F. A. Doria, and J. A. de Barros. A suppes predicate for general relativity and set-theoretically generic spacetimes. *International Journal of Theoretical Physics*, 29(9):935–961, Sep 1990.

[DD12] Francisco Antonio Doria and Manuel Doria. Einstein meets Gödel. *New Mathematics and Natural Computation*, 08(01):123–137, 2012.

[DlCCC+16] G. De las Cuevas, T. S. Cubitt, J. I. Cirac, M. M. Wolf, and D. Pérez-García. Fundamental limitations in the purifications of tensor networks. *Journal of Mathematical Physics*, 57(7), 2016.

[DPR61] Martin Davis, Hilary Putnam, and Julia Robinson. The decision problem for exponential diophantine equations. *Annals of Mathematics*, 74(3):425–436, 1961.

[EMG12] J. Eisert, M. P. Müller, and C. Gogolin. Quantum measurement occurrence is undecidable. *Phys. Rev. Lett.*, 108:260501, Jun 2012.

[Fri21] Tobias Fritz. Quantum logic is undecidable. *Archive for Mathematical Logic*, 60(3):329–341, May 2021.

[FT82] Edward Fredkin and Tommaso Toffoli. Conservative logic. *International Journal of Theoretical Physics*, 21(3):219–253, 1982.

[Göd31] Kurt Gödel. Über formal unentscheidbare sätze der principia mathematica und verwandter systeme i. *Monatshefte für Mathematik und Physik*, 38(1):173–198, Dec 1931.

[GS87] Yuri Gurevich and Saharon Shelah. Expected computation time for hamiltonian path problem. *SIAM Journal on Computing*, 16(3):486–502, 1987.

[Gur91] Yuri Gurevich. Average case completeness. *Journal of Computer and System Sciences*, 42(3):346–398, 1991.

[Hal83] F. D. M. Haldane. Nonlinear field theory of large-spin heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis néel state. *Phys. Rev. Lett.*, 50:1153–1156, Apr 1983.

[JNV+20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*= RE. 2020. arXiv:2001.04383.

[Kan90] I. Kanter. Undecidability principle and the uncertainty principle even for classical systems. *Phys. Rev. Lett.*, 64:332–335, Jan 1990.

[KGE14] M. Kliesch, D. Gross, and J. Eisert. Matrix-product operators and states: Np-hardness and undecidability. *Phys. Rev. Lett.*, 113:160503, Oct 2014.

[Kom64] Arthur Komar. Undecidability of macroscopically distinguishable states in quantum field theory. *Phys. Rev.*, 133:B542–B544, Jan 1964.

[KvdES+22] Andreas Klingler, Mirte van der Eyden, Sebastian Stengele, Tobias Reinhart, and Gemma De las Cuevas. Many bounded versions of undecidable problems are np-hard, 2022. arxiv.2211.13532.

[Lev86] Leonid A Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.

[LWMA17] Miklós Lajkó, Kyle Wamer, Frédéric Mila, and Ian Affleck. Generalization of the haldane conjecture to su(3) chains. *Nuclear Physics B*, 924:508–577, 2017.

[Mat70] Yuri Vladimirovich Matiyasevich. The diophantineness of enumerable sets. In *Doklady Akademii Nauk*, volume 191, pages 279–282. Russian Academy of Sciences, 1970.

[Moo90] Cristopher Moore. Unpredictability and undecidability in dynamical systems. *Phys. Rev. Lett.*, 64:2354–2357, May 1990.

[Moo91] C Moore. Generalized shifts: unpredictability and undecidability in dynamical systems. *Nonlinearity*, 4(2):199, 1991.

[MU16] Alexei Miasnikov and Alexander Ushakov. Generic case completeness. *Journal of Computer and System Sciences*, 82(8):1268–1282, 2016.

[PER81] Marian Boykan Pour-El and Ian Richards. The wave equation with computable initial data such that its unique solution is not computable. *Advances in Mathematics*, 39(3):215 – 239, 1981.

[Ric68] Daniel Richardson. Some undecidable problems involving elementary functions of a real variable. *Journal of Symbolic Logic*, 33(4):514–520, 1968.

[Rob71] Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Inventiones mathematicae*, 12(3):177–209, 1971.

[Sha82] Adi Shamir. A polynomial time algorithm for breaking the basic merkle-hellman cryptosystem. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, pages 145–152, 1982.

[Slo20] William Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 33(1):1–56, 2020.

[SM21] Naoto Shiraishi and Keiji Matsumoto. Undecidability in quantum thermalization. *Nature Communications*, 12(1):5084, Aug 2021.

[SMG$^+$20] G Scarpa, András Molnár, Yimin Ge, Juan José García-Ripoll, Norbert Schuch, David Pérez-García, and S Iblisdir. Projected entangled pair states: Fundamental analytical and numerical limitations. *Physical Review Letters*, 125(21):210504, 2020.

[SS21] Matteo Scandi and Jacopo Surace. Undecidability in resource theory: Can you tell resource theories apart? *Phys. Rev. Lett.*, 127:270501, Dec 2021.

[Tac22] Yuji Tachikawa. Undecidable problems in quantum field theory, 2022. arxiv.2203.16689.

[Tao19] Terence Tao. On the universality of the incompressible Euler equation on compact manifolds, II. Non-rigidity of Euler flows. *arXiv preprint arXiv:1902.06313*, 2019.

[Tur36] Alan Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(1):230–265, 1936.

[VR92] Ramarathnam Venkatesan and Sivaramakrishnan Rajagopalan. Average case intractability of matrix and diophantine problems. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of Computing*, pages 632–642, 1992.

[Wan61] H. Wang. Proving theorems by pattern recognition – ii. *The Bell System Technical Journal*, 40(1):1–41, Jan 1961.

[Wan74] Paul S. Wang. The undecidability of the existence of zeros of real elementary functions. *J. ACM*, 21(4):586–589, October 1974.

[Wan95] Jie Wang. Average-case completeness of a word problem for groups. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '95, page 325–334, New York, NY, USA, 1995. Association for Computing Machinery.

[WB95] J. Wang and J. Belanger. On the NP-Isomorphism Problem with Respect to Random Instances. *Journal of Computer and System Sciences*, 50(1):151–164, 1995.

[WCPG11] Michael M. Wolf, Toby S. Cubitt, and David Perez-Garcia. Are problems in quantum information theory (un)decidable?, 2011. arXiv:1111.5425.

[WOC22] James D. Watson, Emilio Onorati, and Toby S. Cubitt. Uncomputably complex renormalisation group flows. *Nature Communications*, 13(1):7618, Dec 2022. arXiv:2102.05145.

[YA16] Adama Yedidia and Scott Aaronson. A relatively small Turing machine whose behavior is independent of set theory. *Complex Systems*, 25, 2016. arXiv:1605.04343.