games

Article

Nash Equilibria and Undecidability in Generic Physical Interactions—A Free Energy Perspective

Chris Fields [1,*] and James F. Glazebrook [2,3]

[1] Allen Discovery Center, Tufts University, Medford, MA 02155, USA
[2] Department of Mathematics and Computer Science, Eastern Illinois University, Charleston, IL 61920, USA; jfglazebrook@eiu.edu
[3] Adjunct Faculty, Department of Mathematics, University of Illinois at Urbana–Champaign, Urbana, IL 61801, USA
* Correspondence: fieldsres@gmail.com

Abstract: We start from the fundamental premise that any physical interaction can be interpreted as a game. To demonstrate this, we draw upon the free energy principle and the theory of quantum reference frames. In this way, we place the game-theoretic Nash Equilibrium in a new light in so far as the incompleteness and undecidability of the concept, as well as the nature of strategies in general, can be seen as the consequences of certain no-go theorems. We show that games of the generic imitation type follow a circularity of idealization that includes the good regulator theorem, generalized synchrony, and undecidability of the Turing test. We discuss Bayesian games in the light of Bell non-locality and establish the basics of quantum games, which we relate to local operations and classical communication protocols. In this light, we also review the rationality of gaming strategies from the players' point of view.

Keywords: free energy principle; Gödel's theorem; Markov blanket; measurement; Nash equilibrium; quantum reference frame; Turing test; undecidability

check for updates

Citation: Fields, C.; Glazebrook, J.F. Nash Equilibria and Undecidability in Generic Physical Interactions—A Free Energy Perspective. *Games* **2024**, *15*, 30. https://doi.org/10.3390/g15050030

Academic Editor: Ulrich Berger

Received: 31 May 2024
Revised: 21 August 2024
Accepted: 22 August 2024
Published: 26 August 2024

1. Introduction

Since its inception by von Neumann and Morgenstern [1], game theory (GT) has provided a fertile ground for formal studies of algorithmic decidability and undecidability. The definition of equilibrium for non-cooperative games in [1] was restricted to two-person, zero-sum games. J. F. Nash in [2,3] introduced a concept of equilibrium applicable to a much more general class of games, based on best-response strategies, regardless of the number of players and any bounds on the eventual payoff. Specifically, the Nash equilibrium (NE) comprises a set of strategies, one for each of the n game players, with the property that each player's choice of strategy is their best response to the choices of the $n-1$ other players [4]. This principle, as applied to an array of gaming strategies (mixed and pure) of 'best response' determined by probability distributions, profoundly influenced the world of economic game theory, and in recognition of this irreducibly original idea, Nash was awarded the Nobel Prize in 1994 (for the foundations and various interpretations and applications of the theory; see, e.g., [4–6]).

To illustrate the notion of an NE, consider a simple two-player game, the *prisoner's dilemma* (PD). Each round is defined by a payoff matrix in which the best payoff is obtained by defecting (D) when offered cooperation (C), e.g.,

		Bob s_B C	Bob s_B D
Alice s_A	C	$(3,3)$	$(0,5)$
	D	$(5,0)$	$(1,1)$

In such a setting, D is the dominant strategy and (D, D) is the Nash equilibrium. The optimal joint strategy is (C, C), in which both players cooperate, but this is clearly unstable.

While the NE in the PD as defined above is the single point (D, D), redefining the space of possible strategies can render the structure of the NE as an attractor in strategy space arbitrarily complex. We could, for example, allow Bob a continuous distribution of strategies, each labeled with a complex number, and replace the 'C' and 'D' column labels in the table above with '$\{c|c \notin \mathcal{M}\}$' and '$\{c|c \in \mathcal{M}\}$', respectively, with $\mathcal{M}$ the Mandelbrot set, or with $\mathcal{M}$ the pullback attractor of some random dynamical system $\mathcal{D}$ on $\mathbb{C}$ for which such an attractor is defined. While such formal moves are somewhat contrived, the latter may realistically reflect the situation in evolutionary game-theoretic settings [7], where the numbers of distinct actions, and hence, "strategies" available to both an organism or population and its environment (comprising other organisms or populations) are very large, but only difficult-to-define sets of such actions have discernibly distinct short-term payoffs.

A fundamental question is posed by the existence, in all games, of NEs: for a given game $\mathfrak{g}$ and the dynamics ϕ on $\mathfrak{g}$, can it be shown whether ϕ converges to an NE for $\mathfrak{g}$? Various results have been proved, many demonstrating nonconvergence for games with particular structures. Three notable examples from this literature are proofs that (i) game dynamics that are uncoupled, in the sense that each player's choice of strategy depends only on their own payoff function, are generically not Nash-convergent, even for point attractors [8]; (ii) games exist for which all dynamics fail to converge to an NE [9]; and (iii) sufficiently high-dimensional multi-player games can have chaotic attractors, and hence, dynamics that never converge [10], though in some cases, simple heuristics can predict long-term outcomes [11]. The second of these results was proved for degenerate games, i.e., games in which multiple moves have the same payoff; however, the problem of deciding whether a game is degenerate is known to be NP-complete except in special cases [12]. In fact, for finite games the computational complexity of NE has be shown to be that of polynomial parity arguments on directed graphs (PPAD) problems [13]. As informally explained in [14], PPAD is a class of all search problems for which a solution is guaranteed to exist for the same combinatorial reason that every game has at least one NE.

Results such as these are generally obtained by representing games as dynamical systems—hence, the idea of game dynamics—but as might be expected from Gödel's theorem [15], any self-consistent axiomatic system sufficient to represent game dynamics is provably incomplete, in the specific sense that whether a given combination of strategies constitutes an NE is generically unprovable, even for finite games [16]. Gödel undecidability extends beyond the question of NE: the question of strategy convergence in a spatialized prisoner's dilemma (SPD)—effectively, a question of whether a finite, heterogeneous cellular automaton (CA) will evolve into a homogeneous CA given a generic distribution of rules/strategies—has also been shown explicitly to be undecidable [17], as discussed in detail in Section 4.3 below.

While the above kinds of questions arise when games are analyzed as abstract structures, players of games also face decidability questions on each round of play. Is, for example, one's currently cooperative opponent in an iterated PD (IPD) playing tit-for-tat, or is their strategy to defect after n rounds? This question is clearly undecidable at round k for any $k < n$. Indeed if the opponent is modeled as a black box, the undecidability of any question about future strategy is undecidable by Moore's theorem [18], which shows that no finite sample of I/O behavior is sufficient to characterize a generic black box. Turing's imitation game [19], otherwise known as the Turing test, provides a case in point: with the assumption that the interrogator is a Turing machine (TM), the undecidability of whether the respondent is human or machine after any n rounds of questions has been proved explicitly [20].

Here, we study these questions of decidability and convergence to an equilibrium in a setting in which the "game" is a generic physical interaction. This effectively generalizes Milnor's idea of a "game against Nature" [21] to situations in which the payoff to Nature bears no particular relationship to the "player's" payoff. To effect this generalization, we first describe generic physical interactions in terms of "strategies" and the "moves" that they generate. We then follow Feynman [22], and later Friston and colleagues [23–27], in employing the variational free energy (VFE) computed by each of the interacting systems, which treats each system's internal dynamics as a model of the other and measures its prediction error, as an inverse payoff measure. When "computation" is treated simply as a functional interpretation of a physical process [28], this representation is completely generic in both classical [25–27] and quantum [29,30] settings. Having described games in this generic setting, we can appeal to undecidability results or "no-go" theorems provable for generic physical systems to understand why, given that NEs generically exist, failure to converge to an NE can be expected for generic games.

We review the formalism needed to represent generic physical interactions as games in Section 2, and show how normal-form games constitute a special case of such generic games. We then review a number of undecidability results provable in this generic setting in Section 3, and discuss how particular classes of games can be seen as providing a priori answers to otherwise undecidable questions. We interpret Nash's theorem in terms of generic physical equilibria in Section 4, showing that the NE can be identified with instances of generalized predictive synchrony in the classical case and with entangled states in the quantum case. We consider the differences between classical and quantum strategies in Section 5, and also discuss the case in which the players manipulate a shared quantum resource as part of the game. We discuss some remaining issues in Section 6 and conclude with some open questions in Section 7.

2. Representing Generic Interactions as Games

2.1. Physical Interaction Is Information Exchange

Let U be an isolated, finite, physical system. We can, without loss of generality, regard U as comprising m binary degrees of freedom. If we consider these degrees of freedom to be classical bits, the possible states of U are simply the m-bit strings with, e.g., the Hamming distance as a metric; if we consider the degrees of freedom to be quantum bits (qubits), we can represent U by a complex Hilbert space $\mathcal{H}_U$ with dimension $\dim(\mathcal{H}_U) = 2^m$. Here, we employ the quantum formalism for its greater generality. Following the development in [31–33], we select a Hilbert space decomposition $\mathcal{H}_U = \mathcal{H}_S \otimes \mathcal{H}_E$ into a "system" S of interest—which we will regard as the "player"—and its "environment" E which plays the role of "Nature" (E can also be interpreted as a collection of resources, e.g., thermodynamic free energy, stigmergic memory, etc., with some degree of stochasticity). We can then write the interaction between S and E as a Hamiltonian (i.e., total energy) operator $H_{SE} = H_U - (H_S + H_E)$, where H_U, H_S, and H_E are the internal or self-interactions of U, S, and E, respectively. We are interested in the case in which H_{SE} is weak enough that the joint state $|SE\rangle$ (using Dirac's notation) is separable, i.e., not entangled, over the time interval of interest. This allows us to write H_{SE} as

$$H_{SE} = \beta_k k_B\, T_k \sum_i^N M_i^k, \tag{1}$$

where $k =\ S$ or E, k_B is Boltzmann's constant, T_k is temperature, M_i^k are N Hermitian operators with eigenvalues in $\{-1, 1\}$, and $\beta_k \geq \ln 2$ is an inverse measure of k's thermodynamic efficiency that depends on the internal dynamics H_k. This interaction H_{SE} provides a formal description of "playing" the game.

We now assume the *holographic principle* (HP), the claim that no more information can be obtained about a physical system than can be encoded on that system's boundary [34–36]; see [37] for details of how the HP applies in this setting. The HP gives Equation (1) a straightforward topological interpretation. Let $\mathscr{B}$ denote the decompositional boundary given implicitly by the Hilbert space factorization $\mathcal{H}_U = \mathcal{H}_S \otimes \mathcal{H}_E$. Given separability, i.e., $|SE\rangle = |S\rangle|E\rangle$, the entanglement entropy $\mathcal{S}(|SE\rangle)$ across $\mathscr{B}$ is zero. We can, therefore, regard $\mathscr{B}$ as a holographic screen, i.e., an ancillary N-qubit array, separating S from E, and depict H_{SE} as in Figure 1.

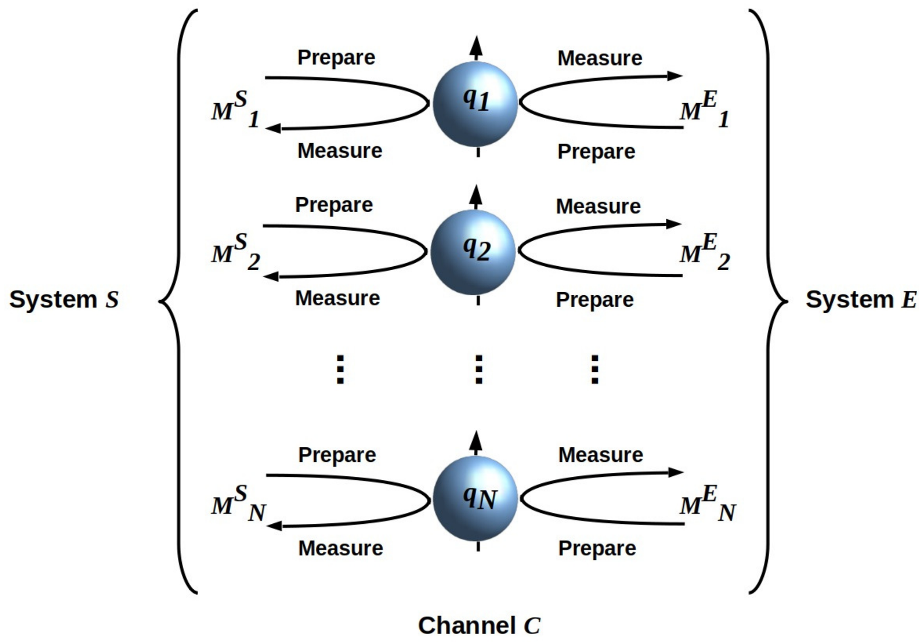

Figure 1. A holographic screen $\mathscr{B}$ separating systems S and E with an interaction H_{SE} given by Equation (1) can be realized by an ancillary array of noninteracting qubits that are alternately prepared by S (E), and then, measured by E (S). Qubits are depicted as Bloch spheres [38]. There is no requirement that S and E share preparation and measurement bases, i.e., quantum reference frames, as described below. Adapted from [33] Figure 1, CC-BY license.

Figure 1 makes explicit a fundamental observation of Wheeler [39]: quantum theory allows any interaction between separable systems to be treated as communication. This communication is bidirectional and indeed informationally symmetric by definition; S and E interact by exchanging N-bit strings. It thus renders the idea of "passive observation" unphysical, as reflected in Wheeler's famous aphorism, "No question? No Answer!".

2.2. Actions Require Quantum Reference Frames

The boundary $\mathscr{B}$ is the "board" (or "media", or "channel") on or through which the game is played. This channel can be any information-encoding space, e.g., the internet in the case of video games, or a private channel in a quantum cryptography setting where an eavesdropper effectively plays a game with a subject who has assumed total privacy [40]. Given $\mathscr{B}$, we can describe the moves and the strategies that drive them. Each move has two components: first S (E) prepares each qubit on $\mathscr{B}$ in some state, after which E (S) measures each qubit. Preparation and measurement, i.e., observation, of qubit q_i are dual processes [41] carried out with the operator M_i^k. This operator is, effectively, an instance of the z-spin operator s_z; it acts on a qubit to prepare it in either the $\uparrow$ $(+1)$ or $\downarrow$ (-1) state. Rendering this action well-defined requires specifying a physical direction that counts as "up" (or $+z$); the opposite direction is then "down" $(-z)$. This specification is achieved by employing a quantum reference frame (QRF), a physical system that provides a fixed, re-usable standard for measurements [42,43]. In a typical laboratory, "up" is defined by the Earth's gravitational field. The QRF is used to "make" the move; choosing a QRF to employ corresponds, in this setting, to choosing a strategy.

When S encodes a bit string on $\mathscr{B}$ by preparing each of the q_i in some particular state, it must select a local $+z_i^S$ QRF for each of the M_i^S, and hence, for each of the q_i. When E then reads a bit string from $\mathscr{B}$, it must also select a local $+z_i^E$ QRF for each of the q_i. If S and E are to remain separable, these choices of local QRFs, which correspond to choices of basis $|i\rangle$ in Equation (1), must be made independently or "freely" [37]; if S's choice of basis depends on E's or vice-versa, they are entangled. If we view $\mathscr{B}$ as a communication channel, independent choice of basis by both S and E can be viewed as introducing noise into the communication; in the extreme case of S choosing $+z_i^S = -(+z_i^E)$ for q_i, S will observe E's encoded bit as being flipped. As U is isolated, there is no classical source of noise in the system; the "noise" due to differences in QRF/basis choice between S and E is purely quantum. The "noise" caused by distinct QRF choices generates statistical surprise, as discussed below; this is the surprise induced when one's opponent chooses a different strategy.

Moves in a game can involve more than one bit; in such cases, a multi-bit strategy is needed. Single-qubit QRFs can be combined to create QRFs that read or write particular bit strings encoded by subsets of qubits on $\mathscr{B}$. We can represent these composite QRFs by hierarchies of maps that form category-theoretic limits and colimits over the relevant single-qubit QRFs [33]; Figure 2 shows such a composite QRF "attached" to $\mathscr{B}$. These *cone–cocone diagrams* (CCCDs) [44,45] are logically regulated, causally and context-sensitive, distributed systems of information flow, as based on the theory of [46]. They are provably general representations of composite QRFs, and are provably equivalent to topological quantum field theories (TQFTs) over the relevant subsets of qubits [47,48]. As with single-qubit QRFs, S and E have independent, free choice of composite QRFs to deploy on $\mathscr{B}$.

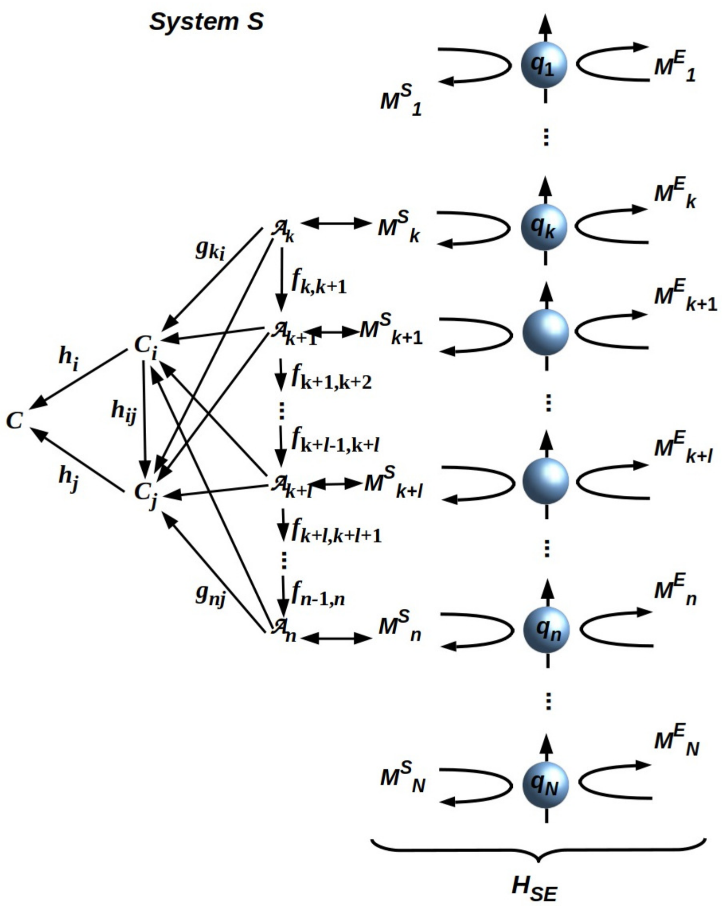

Figure 2. "Attaching" a CCCD to an intersystem boundary $\mathscr{B}$ depicted as an ancillary array of qubits. The operators M_i^k, $k = S$ or E, are single-bit components of the interaction Hamiltonian H_{SE}. The node C is both the limit and the colimit of the nodes $\mathcal{A}_i$; only leftward-going (cocone-implementing) arrows are shown for simplicity. See [29–31,47] for details. Adapted from [31], CC-BY license.

We can now fully describe a move in a generic S-E game. Assuming S has the first move, S deploys some QRF/strategy Q_i^S to encode a particular bit string on the subset $\mathrm{dom}(Q_i^S)$ of qubits on $\mathscr{B}$, after which E deploys some QRF/strategy Q_j^E to read a bit string from the subset $\mathrm{dom}(\mathrm{Q}_j^E)$ of qubits. The turn then reverses, with E encoding and S reading. Note that nothing requires that $\mathrm{dom}(\mathrm{Q}_i^S) = \mathrm{dom}(\mathrm{Q}_j^E)$, and nothing requires that either S or E deploys the same QRF/strategy on each move. In a generic game, both players can be expected to deploy multiple QRFs/strategies, up to some limit imposed by their available computational resources.

2.3. VFE Provides a Generic Payoff Function

From a global perspective, the alternating moves of the S-E game are driven by the global self-interaction H_U; interposing $\mathscr{B}$ between S and E does not affect this global interaction in any way. From the perspective of S or E, the moves are driven by the data that the other party encodes on $\mathscr{B}$. We can also say: they are driven by how the internal dynamics H_S and H_E respond to the perturbations of the system states $|S\rangle$ and $|E\rangle$, respectively, by the interaction H_{SE}.

The *free energy principle* (FEP), introduced by Friston and colleagues [23–27], provides a statistical physics representation of the above facts. Informally, the FEP states that S and E will remain distinct only if they remain sufficiently sparsely or weakly coupled that the boundary between them remains well-defined[25]. If U is treated as a classical causal network, the boundary becomes a *Markov blanket* (MB), as originally defined by Pearl [49]. In this setting, the FEP can be formulated as the requirement that states of S and E each remain in the vicinity of some respective *non-equilibrium steady state* (NESS) [25], or that almost all paths through the joint space that begin in $S(E)$ remain in $S(E)$ [27]. We can, clearly, re-interpret these classical statements simply as requiring that S and E remain unentangled, i.e., that both have "internal states" that contribute only negligibly to H_{SE}.

The utility of the FEP as a guiding principle is that it shifts the focus from characterizing the internal dynamics H_S or H_E, to characterizing the function that each must perform to maintain the long-term integrity of its MB, i.e., of $\mathscr{B}$. Again speaking informally, S's ability to maintain a well-defined boundary—if S is an organism, to stay alive—depends on keeping environmental perturbations of its state relatively small. This can be formulated in terms of *prediction* and *surprise*: whatever H_S does, it needs to minimize the surprise $-\ln p(b)$, where b is an MB or boundary state (in the notation of Section 2.1, of the holographic screen $\mathscr{B}$) relative to a prediction η of E's behavior. The *variational free energy* (VFE) measured at $\mathscr{B}$ is an upper bound on surprise ([25] Equation (2.3)):

$$\begin{aligned} F &= D_{KL}[q_\mu(\eta)|p(\eta)] - \mathbb{E}_q[\ln p(b|\eta)], \\ &= D_{KL}[q_\mu(\eta)|p(\eta|b)] - \ln p(b), \end{aligned} \tag{2}$$

where $q_\mu(\eta)$ is a variational density over predicted external states η parameterized by internal states μ, and $\mathbb{E}_q$ is an expectation value operator parameterized by the variational density q. Note that the Kullback–Leibler (KL) divergence in the second equality scores the (non-negative) prediction error as a divergence between the variational prediction and the true distribution over external states, given observable MB/boundary states. Because this prediction error is non-negative, the VFE furnishes a bound on surprise, which becomes exact when the prediction error is zero.

The first equality in Equation (2) expresses VFE in terms of complexity minus accuracy (first and second terms, respectively), where there is an intimate relationship between the complexity (i.e., divergence between the variational posterior and prior) and the algorithmic complexity of the generative model as a description of E. Heuristically, this means that minimizing VFE provides the simplest accurate account of an environment that can never be observed directly. In turn, minimizing complexity speaks to a description of the environment in terms of minimum message or description lengths, i.e., that speaks to universal computation [50–53].

We can now state the FEP as the claim that S and E will remain distinct systems to the extent that their respective dynamics H_S and H_E are able to each minimize the VFE F measured at their side of $\mathscr{B}$, i.e., for their own inputs and predictions. Note that as they are described by Equation (2), S and E are equally "in the game" of maintaining their mutual distinction.

In the Bayesian sense, μ can be seen as encoding a posterior over the external state η. Minimizing the VFE leads to minimizing a prediction error, a process encompassed by a *generative model* (GM) implemented by the agent's internal dynamics. Here, it is apt to see VFE minimization, from an informational perspective, as a thermodynamically driven process, assimilating the dynamics of an environment whose thermodynamic agency becomes minimized while driving internal self-organization. Declining VFE is then *self-evidencing* in the literal sense of providing evidence for the implementing system's continuing existence [25]. From a dynamical systems perspective, this mechanism maintains the internal state μ in the neighborhood of an NESS solution to the system's density dynamics as given above.

Nothing in the above assumes anything particular about H_S or H_E beyond the function of maintaining their distinctness, or in quantum language, the separability of S and E. *The FEP therefore describes generic interactions as a simple game with VFE minimization as the payoff function.* The objective of the game is maintaining a distinct existence with self-evidencing. It is the most basic game any system plays, and all systems play it all the time.

2.4. Normal-Form Games Are Special Cases

It is implicit in game theory that the players exist and are distinct from one another; normal-form games are, therefore, special cases constructed on top of the generic "game of maintaining existence" described above. Most games are not, moreover, games against all of Nature, i.e., all of E, but games against some components of E, with the other components being neutral, providing infrastructure, or simply being neglected altogether.

The QRF formalism illustrated in Figure 2 allows us to say precisely what is meant by S identifying and interacting specifically with some "system" X embedded in E. To remain identifiable by S over time, X must have some component X_R with a state $|X_R\rangle$ (or state density ρ_{X_R}) that is invariant under the interaction H_{SE}. To accomplish the identification of X over time, S must implement some QRF $\mathbf{X_R}$ specific for X_R, i.e., that produces an outcome '+1' when $|X_R\rangle$ is detected and '-1' otherwise. To be seen by S as making "moves" of interest, X must have some component X_P (a "pointer" component) with a state $|X_P\rangle$ that varies under the interaction H_{SE}, and S must implement a QRF $\mathbf{X_P}$ that specifically detects this variable "pointer state" [29,31,33]. The sector $\mathrm{dom}(\mathbf{X_R}) \cup \mathrm{dom}(\mathbf{X_P})$ of $\mathscr{B}$ is, effectively, the "image" of, or in the terminology of [54] the "icon" of, X for S.

Playing a multi-move game with X requires that S has some memory for previous moves. We could consider S to have a purely procedural memory, e.g., to implement some learning algorithm that updated its decision algorithms on every cycle, as is standard for artificial neural networks [55]. More interesting from the present perspective is the case in which S has declarative memories of particular events, and so can trace X's behavior explicitly through time. Writing and then reading a declarative memory Y requires a dedicated QRF $\mathbf{Y}$, as shown in Figure 3. The process of irreversibly writing a declarative memory requires both thermodynamic energy from the environment and an internally counted time, which we can represent by a counter, or time QRF, $\mathcal{G}_{ij}$ [31,33].

In order to participate in a two-player game with X, therefore, S needs a QRF $\mathbf{X}$; two memories $\mathbf{Y}_S$ and $\mathbf{Y}_X$ of self- and X-actions, respectively, both of some temporal depth $\Delta t_S \geq 1$; a VFE measure F_X over the sector $\mathrm{dom}(\mathbf{X})$; and a decision function D that computes what to do in the next timestep. Generalizing to k players is formally straightforward. We can see in this a generalization of normal form, which replaces the $k^2 - k$ interplayer QRFs with k "objective" players, the decision functions D_i with sets of discrete strategies, and the VFE measures F_{ij} with maps from selected actions to $\mathbb{R}$. We can, in other words, see normal form as an assumption of both classical realism—the players are assumed to perceive and act within an observer-independent "game world"—and discreteness—sets of strategies and associated payoffs are assumed to be finite, or even tractably small. These conditions formalize, as they were intended to, our intuitive notions of what a "game" is.

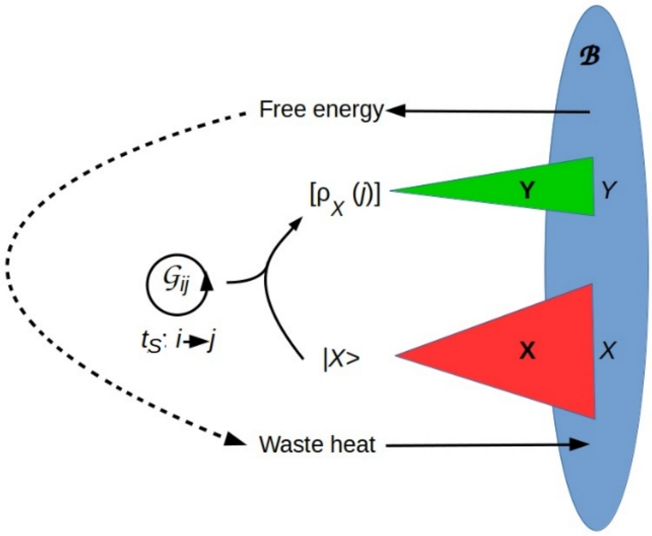

Figure 3. Cartoon representation of a system A that deploys a QRF **X** (red triangle) to measure the state of an external system X in its informational environment (i.e., a sector X of its boundary $\mathscr{B}$), and then, deploys a second QRF **Y** (green triangle) to write the outcome to a memory sector Y. This process induces one "tick" of an internal clock $\mathcal{G}_{ij}$ that defines an internal elapsed time t_S. The process is powered by a thermodynamic loop from (thermodynamic free energy in) and back to (waste heat out) the physical environment E. Adapted with permission from [37], CC-BY license.

2.5. Example: The IPD as a Prediction Game

The IPD again provides a simple example that illustrates how normal form abstracts from the physical description. The players—Alice and Bob—are embedded to some overall shared environment that provides them with resources, particularly thermodynamic free energy, that allow them to play the game iteratively. They are assumed to have identified each other as players, and to each be able to focus their attention exclusively on the other. The VFE for each player, in other words, is assumed to be a function only of what the other player does; sources of uncertainty in the general environment are viewed as negligible or irrelevant. Clearly this is an abstraction of any real setting, e.g., the real situation of any pair of organisms. The players are also each restricted to one or the other of the same two possible moves on each cycle, again a considerable idealization of most real settings.

Each player in the IPD has a model of the other, or more specifically, a move-to-move updated probability distribution over the other's next moves. The IPD is a dilemma because these probability distributions are typically not unimodal, and even if they are unimodal—corresponding to a "certain" prediction—they may not be predictively accurate. Decoupling between subjective probabilities and actual outcomes is, of course, typical in real situations.

The moves in the IPD can be viewed as predictions: a C move predicts a C on the opponent's part, while D predicts D or, more hopefully, C. Observing D after C or C after D are both surprising, but in opposite directions: D after C is disappointing and decrements the model probability that the opponent is a cooperator, while C after D increases that probability. Predicting good results from risky moves is a key component of active exploration of the environment driven by the FEP, often called "epistemic foraging" [25]. While in the abstracted context the IPD it can appear deceitful, in other situations such high-risk/high-reward behavior is a key indicator—and result—of intrinsic motivation [56]. It is, for example, a central driver of science [57].

The IPD converges to (D,D) when the players can, effectively, no longer learn any more about each other. This is an example of generalized synchrony, the generic equilibrium state for systems interacting via the FEP discussed in Section 4 below. While this NE exists, when it will be reached is unpredictable in real time by the players, as further discussed below.

The largest abstraction from reality made in the PD, and hence, in the IPD, is the assumption that the game is not uncoupled, i.e., that the players each know the entire payoff matrix. This information is provided a priori; there is no "bidding phase" in a PD where the players optionally reveal information about their payoffs and bluffing is allowed, so that any resulting "knowledge" of other players' payoffs, even if subjectively certain, may be highly inaccurate. As shown in [8], uncoupling generically disrupts convergence, even to unique point NEs. The next section discusses this relationship between access to information and convergence to equilibrium in generic physical terms.

3. Generic Limits on Observation

The generic description of physical interactions as information exchange sketched above allows us to prove a number of no-go results that place generic limits on what can be deduced from, or decided on the basis of, finite observations. These results apply to all players of all games. They can be viewed as extensions or generalizations of Moore's theorem for observations of classical black boxes [18]. Defining the dimension $\dim(Q)$ of a QRF Q as 2^j, where j is the number of binary degrees of freedom required to implement Q, we proved in [58] that:

Theorem 1 ([58] Theorem 1)**.** *Let S be a finite system and Q be a QRF implemented by H_S. The following statements hold:*

1. *S cannot determine, by means of Q, either Q's dimension $\dim(Q)$, Q's associated sector dimension $\dim(\mathrm{dom}(Q))$, or Q's complete I/O function.*
2. *S cannot determine, by means of Q, the dimension, associated sector dimension, or I/O function of any other QRF Q' implemented by S.*
3. *S cannot determine, by means of Q, the I/O function or dimension of any QRF Q' implemented by any other system S', regardless of the relation of S to S', from $S' = S$ to $S' = E$, inclusive.*
4. *Let $S = S_i S_j$, in which case $E_i = ES_j$. Then, S_i cannot determine, by means of a QRF Q_i, the I/O function or dimension of any QRF Q_j implemented by S_j.*

Proof. See [58]. All clauses follow from the inability to specify H_S or H_E given only H_{SE}, or in particular, the finite set of bits encoded on some observable sector of $\mathscr{B}$. □

A corollary of Theorem 1 is that GMs implemented by physical systems are inevitably incomplete, in the sense that there are inputs that can be received but not predicted, and adding more or different QRFs or hierarchical (i.e., meta-) layers cannot make them complete ([58], Corollary 3). This sense of incompleteness is obviously reminiscent of Gödel's theorem, and follows from Gödel's theorem immediately if the GMs in question are treated as axiomatic systems with at least the power of arithmetic.

A further result that will be relevant to Section 5 below is that no system can determine the entanglement entropy across its own boundary ([59], Corollary 3.1). This restricts any system from determining, by observation, that it is not entangled with its environment, and hence, that it has no "back channel" of communication with its environment that is not mediated by classical information. The result follows from the undecidability of the problem of determining whether an action alters the entanglement entropy of the environment [59]. This is a quantum version of the classical frame problem, the problem of deciding what remains invariant after an action [60].

Theorem 1 implies that a system cannot determine its own decision function D, and hence, that it cannot report its decision function to any other system. It moreover implies that no system can reliably infer the decision function of any other system by observation. Hence, no player of a generic game can report its own strategy to other players, or reliably infer the strategies of other players from their behavior. Theorem 1 likewise restricts any system from reliably inferring the VFE measurements, and hence, the payoff function, of other systems. Generic games are, therefore, all uncoupled in the sense of [8], and hence, are incomplete-information, or Bayesian, games as defined by Harsanyi [61]. This uncoupling extends to the temporal depth or reliability of other systems' memories for previous moves. Theorem 1 thus shows that games with well-defined rules that are known by all players effectively postulate shared a priori knowledge that cannot be obtained empirically. In practice, we can think of this as an assumption of a shared language, or a shared semantics. Such assumptions are known to be problematic on classical, logical grounds [62]; we will see this more deeply when considering communication protocols explicitly in Section 5 below.

4. Convergence and Equilibria in Generic Interactions

4.1. Convergence Driven by the FEP

Let us now consider what happens when two agents, both driven by the FEP, each strive to reduce the VFE they measure on their respective sides of their mutual boundary. For each agent, reducing VFE is increasing the accuracy of their GM for predicting the other agent's behavior. Each can, therefore, be expected to engage in a combination of learning any patterns in their opponent's behavior and acting on their opponent in order to either alter their behavior or induce new learnable patterns. The combination of these VFE-reduction strategies is *active inference* in the language of Friston and colleagues [24,63–65]. As active inference minimizes VFE, not surprise per se, it can be viewed as approximate Bayesian inference.

Players of games against Nature frequently lose; Nature has many possible moves and is notorious for radically and capriciously changing the rules [66]. For the present purposes, however, we are mainly interested in active-inference games in which H_{SE}, and hence, the "rules" remain fixed, allowing the players to decrease their VFE by improving their predictions for a finite, but unbounded, number of rounds. This scenario naturally raises three questions:

- Does an equilibrium that minimizes VFE—maximizes predictive accuracy—for both S and E exist?
- Can S or E determine by finite observation that they have reached such an equilibrium?
- Can S or E determine that they are on a trajectory toward such an equilibrium?

Nash's theorem [3] gives a positive answer to the first of these questions. Such equilibria have also been characterized in the classical FEP literature, where it has been shown how mutual predictability induces generalized synchrony [67,68]; see, e.g., [69–72] for relevant earlier work and [73,74] for applications. A simple and limiting example of generalized synchrony is convergence to thermodynamic equilibrium, in which prediction ceases at "stasis" because there is no longer available thermodynamic free energy to support computation. Unlike the classical formulation, the quantum formulation of the FEP does not employ an embedding space to enforce separability between S and E; here, the limit of perfect mutual predictability corresponds to entanglement [29], as discussed further in Section 5 below.

As could be expected from the work of Hart and Mas-Colell on uncoupled games [8], the answers to the second and third questions above are negative. As noted earlier, this is evident even in as simple a game as an IPD: no sequence of cooperate (C) moves is sufficient to rule out the next move being defect (D). Even well-supported predictions, in other words, can fail. One reason for this follows immediately from the opacity of MBs discussed in Section 3 above: S and E cannot determine by observation how much internal memory their opponent has, so cannot predict with reliability their opponent's planning horizon.

A deeper reason for the inability of players to determine whether they are on a path to convergence is provided by Rice's theorem [75]. No restrictions have been placed on the internal dynamics H_S or H_E; hence, they can implement arbitrarily complex programs. Rice showed that no TM can determine the function computed by any arbitrary program. Hence, S and E could not determine each other's strategies, in the general case, even if they had full access to a program describing their dynamics. Indeed, they cannot determine whether such a program—and hence, the represented dynamics—halt on a given input [76,77].

Yet another perspective on convergence is provided by the undecidability of the classical frame problem [78]. This prevents S or E from reliably predicting the result of a perturbative action, and hence, from reliably predicting what a given move will reveal about their opponent's strategy. A final perspective is provided by Gödel's theorem itself: there are attractors that cannot be proven to either be or not be NEs [8,16].

4.2. Example: IPDs and Generalized Imitation Games

Turing's imitation game [19] is introduced via the question 'can machines think?' and is typically regarded in this context. It can, however, also be seen as an exemplar of a particular kind of active-inference game, one in which the "opponent's" objective is to cause the "player" to build a GM of the opponent's behavior that is false. In Turing's original version, the interrogator must decide which of a computer and a human is which, while both attempt to deceive them. We will refer to all games of this sort as *generalized imitation games* (GIGs).

Sato and Ikegami showed Turing's imitation game is Turing-undecidable after any finite number of rounds. As noted earlier, we can see this result as a special case of Moore's theorem [18]: finite observations are, in general, insufficient to reveal the dynamics unfolding inside an MB. Moore's theorem applies to any GIG in which the players are separable; hence GIGs in general are undecidable after any finite number of rounds.

Strategic deception is a central feature of human social behavior [79]; hence, many social interactions are, at least in part, GIGs. Successful deception requires multi-timestep memory, both in deceiver and deceived; all GIGs are, therefore, multi-round games. An explicit, move-by-move record, in particular, serves as an additional testing resource for GMs that are evolved on each timestep and encode prior behavior only implicitly. Hence, skilled players of GIGs are "fast learners" who also have "good memory".

As pointed out in [80] as a reflection on [1], "conscious choice" is not assumed in GT, which treats games as formal structures and the agents that play them as, effectively, algorithmic systems (see also a similar discussion in [81], and see [82] for a review of empirical evidence that "conscious choice" is an illusion even in awake humans who explicitly report it). Notions of learning, memory, strategy, and deception are interpreted along these same lines. While it is commonplace to treat GT agents as TMs, a resource theory perspective that places specific limitations on memory or processing capability can yield useful insights (see also, e.g., [83] for the relationship of the logical studies of Gödel, Turing, and Post to the question of how digital agents can innovate, and [84] for a survey of ideas of the former).

Consider, once again, an IPD, which can be thought of as a GIG in which players can surprise each other by shifting strategies. As pointed out earlier, away from the (D,D) equilibrium, not even a complete, explicit record of all past moves is sufficient to reliably predict one's opponent's next move. Taiji and Ikegami [85] have studied IPDs from a resource theoretic perspective, implementing each player as a recurrent neural network (RNN) to enforce a purely implicit memory of previous moves. If identical learning algorithms are employed for each player, the "Bayesian prior" that distinguishes them becomes the algorithm that employs the memory encoded by the RNN to choose a next action. Taiji and Ikegami studied two such priors, "pure reductionist Bob", that employs the RNN as a model of Alice's past moves; and "clever Alice", that employs the RNN to simulate pure reductionist Bob. Effectively, clever Alice builds a model of herself—the RNN model that pure reductionist Bob would build—and treats this model as the GM employed by her opponent. This attempted 'mirroring' of each player's strategies and predictions in the IPD has also been studied as a type of imitation game in [86] (see also the discussion in [85]). As might be expected, IPDs that pit these players against each other, or against copies of themselves, eventually converge to the (D, D) equilibrium. As pointed out in [85], the (C, C) strategy is stable only if each player has an essentially unshakable prior that the opponent plays tit-for-tat. This outcome can be overcome if the players adopt a quantum strategy [80,87], the basic elements of which we will review in Section 5 below.

There are a number of slants on the IPD. Of these, consider the following example. Suppose we take a large number N of rounds of play, whereby cooperation (C) is seen as the best long-term payoff. The snag is, however, that a player might defect (D) at the final round with nothing to lose, in which case the opponent has no opportunity to react. But if defection for both is known for round N, then the same scenario arises, without loss, in round $N - 1$, and so on [14]. But [14] notes, following [88], that this apparent anomaly can be resolved if the two players each have sufficiently small memories, which assumes they are finite automata with k states, $2 \leq k < N$. Then, cooperation can be restored as an equilibrium, in contrast to the previous scenarios, since given the lack of memory to count up to N, as known to both players, any intermediate strategies as evoked in the previous case will no longer apply.

We can view these as carefully contrived applications of the *good regulator theorem* of [89]: any good regulator, i.e., one that is maximally successful and simple, must be an isomorphic model of the system being regulated. The results of Section 3 show that no agent can reliably infer that it is such a good regulator. The infeasibility of isomorphic regulators in practical applications is well known; practical regulators are coarse-grained models, not isomorphic models. The FEP builds this in by employing the variational approximation of Equation (2); see [90] for a recent discussion.

4.3. Example: The SPD and Its Limit Sets

The *spatialized prisoner's dilemma* (SPD) adds a spatial dimension to the time dimension of the IPD. In the SPD, players compete against their eight nearest neighbors, replicating the winning strategy after each full round if defeated by any neighbor. Players are, effectively, cells of a finite CA embedded in some infinite background, each of which implements the same updating rule but whose game strategies can differ. Grim [17] has shown that as finite configurations of IPD strategies, SPDs are formally undecidable; specifically, that *for any chosen infinite background within an SPD, there is no algorithm which reveals in every case whether or not an embedded finite array of strategies will result in a progressive conquest by a single strategy.* Stable patterns obtained after multiple SPD rounds, whether these are uniform, and thus, correspond to "conquest by a single strategy" or not, are limit sets of the CA. We can, therefore, see the undecidability of SPDs as a special case of Kari's theorem [91], which states that all nontrivial properties of such limit sets are undecidable. This result has been extended to show that some properties of CA dynamics, including whether the time evolution map is the identity map, are undecidable even when limited to evolution within a limit set [92].

The SPD imposes, in effect, a non-local prediction problem on its players: an opponent's changes in strategy in the SPD are not capricious, but are rather determined by the behavior of neighbors that are near the opponent but distant to the predicting cell. It can thus be considered a "motivated" GIG, where the motivation comes from outside the immediate two-player setting. Undecidability at the level of the players—by analogy to [20]—derives in this case not from indeterminism, but from lack of access to distantly acting information.

4.4. Prediction, Regulation, and Generalized Synchronization—A Circularity of Idealizations

We see in the examples of IPDs and SPDs that the idea of "deception" employed to define GIGs as a class is just a convenient shorthand for what time and space render inevitable—the impossibility of reliably predicting what will happen next. Adopting a GT perspective on generic physical interactions—or on behavior driven by the FEP—thus emphasizes how apparent minima of VFE can become unstable, regulators that appeared good can be thrown out of their operating windows, and states of generalized synchrony can collapse back into chaos. In a game against Nature, in particular, the limit of identical synchronization, in which both state spaces and GMs are isomorphic, is obtained only when Nature is cut precisely into identical halves. We arrive, then, at an unfolding circularity of idealized concepts:

Good Regulator $\Longrightarrow$ *Identical Synchronization* $\Longrightarrow$ *Winnable Imitation Game (or decidable Turing Test)* $\Longrightarrow$ *Good Regulator*

The source of undecidability in every case is clear: if two interacting systems are identical, self-reference and other-reference are indistinguishable, and Gödel's theorem applies equally to either.

As noted earlier, the "good regulator" limit in the quantum formulation of the FEP is entanglement. We now turn to games that employ this non-classical resource.

5. Quantum Games

5.1. Definitions and Formalism

The formalism for generic interactions in Section 2 applies to quantum as well as classical systems, and hence, can be applied to games involving quantum interactions. Since the pioneering work of Meyer [93], specifying two-person zero-sum games in terms of a notion of quantum strategy and (mixed) quantum equilibria, the interplay between GT and quantum information has produced a quantum theory of games (see, e.g., the review of [94]). This generalizes classical GT by involving, for instance, quantum (mixed) strategies, non-locality, entanglement, and other quantum concepts, and entails far-reaching consequences beyond the classical case. The most notable of these is that quantum strategies are superior to classical strategies, and thus, the expected payoffs can be considerably greater. In short, any quantum system which can be manipulated by two or more parties, in which the utility of the moves can be suitably quantified, qualifies as a quantum game. In terms familiar from quantum computing and quantum information, these include optimal quantum cloning, eavesdropping in quantum cryptography (see listings in [80,87], and, e.g., [40,95]), quantum entanglement and evolution of cellular automata [96], quantum entanglement and secret sharing [97]. The most general form of the classical PD can, moreover, be faithfully represented in the quantum PD [87]. All of these are, as we will see, instances of *local operations, classical communication* (LOCC) protocols [98], in which two or more agents manipulate some resource while also exchanging information about their activities via some separate, classical communication channel.

To an extent, our preliminary account follows [80,87], which is broadly applicable to the basics of the subject as developed by other authors (e.g., [93]), and for which the classical version of the game is often faithfully represented in the quantum version. The key differences exhibited by the latter involve linear superposition of actions/strategies, entanglement between the players, and adoption of quantum probabilities. Let us proceed with a formal description:

Definition 1. *A two-player quantum game* $\Gamma = (\mathcal{H}, \rho, S_A, S_B, P_A, P_B)$ *consists of:*

(a) *The underlying Hilbert space* $\mathcal{H}$ *of the physical system;*
(b) *The initial state* $\nu \in \mathcal{S}(\mathcal{H})$*, where* $\mathcal{S}(\mathcal{H})$ *is the associated game-state space;*
(c) S_A *and* S_B *are sets of permissible quantum operations of players A and B, respectively;*
(d) P_A *are* P_B *are the utility functions specifying the respective utility for each player.*

A *quantum strategy* $s_A \in S_A, s_B \in S_B$ is a quantum operation; in formal terms, a completely positive trace-preserving self-map $\mathcal{S}(\mathcal{H}) \longrightarrow \mathcal{S}(\mathcal{H})$. Quantum games also include various implicit rules depending on the nature of the game in question. A quantum game is a *zero-sum game* if the expected payoffs sum to zero for all pairs of strategies; that is, when $P_A(s_A, s_B) = -P_B(s_A, s_B)$. Otherwise, it is a*non-zero-sum game*.

Next, we consider how equivalence between strategies is defined. Suppose Alice has access to two quantum strategies s_A and s'_A. They are said to be *equivalent* if $P_A(s_A, s_B) = P_A(s'_A, s_B)$, and $P_B(s_A, s_B) = P_B(s'_A, s_B)$, for all of Bob's possible strategies s_B. In other words, s_A and s'_A yield the same expected payoff for both players, for all possible s_B. Likewise, equivalence of strategies for Bob is defined. There are several analogous concepts in quantum GT extending those in classical GT which we will itemize below (again referring to [80,87]):

Definition 2. *(a)* *A quantum strategy is called a dominant quantum strategy of Alice if $P_A(s_A, s'_B) \geq P_A(s'_A, s'_B)$ for all $s'_A \in S_A, s_B \in S_B$. Likewise, a dominant strategy for Bob is defined.*

(b) *A pair (s_A, s_B) is said to be an equilibrium in dominant strategies if s_A and s_B are the players' respective dominant strategies.*

(c) *A combination of strategies (s_A, s_B) is called a Nash equilibrium if*

$$P_A(s_a, s_b) \geq P_A(s'_a, s_B)$$

$$P_B(s_A, s_B) \geq P_B(s_A, s'_B)$$

for all $s'_A \in S_A$ and $s'_B \in S_B$.

(d) *A pair of strategies (s_A, s_B) is called Pareto optimal if it is not possible to increase one player's payoff without reducing the other's payoff.*

The single-shot PD provides an example. The quantum version of the game commences by assigning to each player the classical strategies C and D, corresponding to two basis vectors, $|C\rangle$ and $|D\rangle$, in the Hilbert space of a two-state system, a quantum bit or qubit. At each instance of the game, the state of play is described by a vector in the tensor product space spanned by the classical basis $|CC\rangle, |CD\rangle, |DC\rangle, |DD\rangle$ (according to the qubit of Alice in the first place, and that of Bob second). More formally, the quantum version of this classical binary choice game arises by the preparation of two qubits by some arbiter who relays these to A and B, who have the necessary devices at hand to manipulate their qubits effectively. Then, they eventually relay these back to the arbiter who implements a measurement to determine the payoff. Note that "manipulation" here can, in principle, result in any arbitrary superposition $\psi = \alpha|C\rangle + \beta|D\rangle$, with $\alpha, \beta \in \mathbb{C}$, and that Alice and Bob must agree, via classical communication, to use the basis vectors $|C\rangle$ and $|D\rangle$, and hence, must agree to deploy QRFs specifying those basis vectors.

This formal description entails a quantum system with underlying Hilbert space given as a tensor product $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, where $\mathcal{H}_A = \mathcal{H}_B \cong \mathbb{C}^2$, with associated state space $\mathcal{S}(\mathcal{H})$. Quantum strategies $s_A \otimes 1_B$ and $1_A \otimes s_B$ are identified with s_A and s_B, respectively, while A and B are disposed to choosing any quantum strategy in the set S of these of which they are aware, but are in principle unaware of whatever strategy each other will take. Applying quantum strategies leads to a map:

$$s_A \otimes s_B : \mathcal{S}(\mathcal{H}) \longrightarrow \mathcal{S}(\mathcal{H}) \tag{3}$$

and from the initial state ν, the system moves to a final state:

$$\sigma = (s_A \otimes s_B)(\nu) \tag{4}$$

When s_A and s_B are unitary operations, then they can be identified with unitary operators U_A and U_B, and can be expressed by $s_A \sim U_A$ and $s_B \sim U_B$, respectively. Hence, the final state σ above can be written as

$$\sigma = (U_A \otimes U_B)(\nu)(U_A \otimes U_B)^\dagger \tag{5}$$

The operations s_A and s_B do not, however, have to be unitary, but could involve measurement, and hence, projection onto some basis vector. Additional technical details of how a quantum game $\Gamma = (\mathbb{C}^2 \otimes \mathbb{C}^2, \nu, S_A, S_B, P_A, P_B)$ unfolds in terms of strategies and payoffs are given in [80,87].

5.2. Example: The Decoherence Game

An isolated quantum system remains in a coherent or "pure" state, e.g., an isolated qubit remains in a state describable as $\psi = \alpha|\uparrow\rangle + \beta|\downarrow\rangle$ in an arbitrary basis $(\uparrow, \downarrow)$. When such a state is exposed to some environment E, it "decoheres" by losing coherence to

E [99,100]; see [101] for a comprehensive review. We can represent this process as a game similar to the quantum PD sketched above, in which quantum states of the qubit S of interest are the "players" and its environment E is the "arbiter." For simplicity, we will just consider two players, which we can represent as the states $|\uparrow\rangle$ and $|\rightarrow\rangle$, where $|\rightarrow\rangle = (1/\sqrt{2})(|\uparrow\rangle + |\downarrow\rangle)$.

When the game begins, S is in a coherent state, so $|\uparrow\rangle$ and $|\rightarrow\rangle$ have equal probabilities. We then introduce interaction with E, which we can represent by Equation (1). Let this interaction begin by being very weak and slowly increase in strength; we can think of T_E in Equation (1) increasing slowly, or of the frequency with which E interacts with S increasing slowly. As required by Equation (1), E chooses a basis for the operator M^E (we need consider only one); we will assume E chooses $(\uparrow, \downarrow)$. This amounts to E choosing the payoff matrix: with the choice $(\uparrow, \downarrow)$, $|\uparrow\rangle$ receives one "outcome point" (i.e., the outcome '1' is indicated by a "detector click") while $|\rightarrow\rangle$ receives zero (no click). When the game begins, $|\uparrow\rangle$ and $|\rightarrow\rangle$ have equal scores, but as H_{SE} gains strength, $|\uparrow\rangle$'s score slowly increases while $|\rightarrow\rangle$'s score does not; if we read probability as proportional to score and normalize, $|\uparrow\rangle$'s probability increases while $|\rightarrow\rangle$'s decreases. Asymptotically, $|\uparrow\rangle$ has probability one and $|\rightarrow\rangle$ has probability zero.

This process of measurement by E "rewarding" the state of S that is an eigenstate of E's chosen basis is called "einselection" [101], projective measurement, or the "collapse of the wavefunction." When E is taken to be large enough to interact with many independent observers—e.g., when E is the ambient photon field—E can be regarded as "encoding" the selected state of S with sufficient redundancy that all observers can detect it; this competitive state encoding is called "quantum Darwinism" and is professed to explain the emergence of a "public" quantum-to-classical transition [102,103]. Like other quantum games, quantum Darwinism is an LOCC protocol: the multiple observers interact with E as a quantum resource while agreeing classically to employ the same measurement basis, e.g., to employ vision and a common conception of what counts as an "object" [48].

5.3. Example: The Bell/EPR Game

Bell/EPR experiments are the "gold standard" for detecting quantum entanglement, and provide the conceptual basis for all other quantum communication protocols [104,105]. In the experiment's canonical form, a centrally located source distributes an entangled two-qubit state—e.g., an entangled photon pair—to two observers, Alice and Bob, who are located at equal distances from the source, but in opposite directions. Each observer is equipped with a spin-orientation detector—e.g., a polarizing filter—that can be set to any direction. Alice and Bob know the frequency with which the source emits entangled states, and independently set their detection directions during the time required for entangled states to reach their locations from the detector; their mutual separation is chosen to both allow this to happen and to prevent collusion between them. They each record their separate observations of each state, and later exchange their results, via a classical channel, in order to compute the statistics of their joint observations. A violation of Bell's inequality [106,107] indicates detection of entanglement; see, e.g., [108] for an informal discussion of both the experiment and the statistical analysis.

A Bell/EPR experiment can be viewed as a two-player, limited-cooperation game against Nature, which both supplies the entangled states and rewards detection of entanglement [109]. The players, Alice and Bob, must agree to make spin measurements on every round, but are forbidden from sharing information about their chosen measurement directions. The score is the cumulative value of a joint-measurement statistic, e.g., the Clauser–Horne–Shimony–Holt (CHSH) statistic [110]. Alice and Bob can maximize their score by making their measurements 45 deg apart, up to the limit of Tsirelson's bound, $2\sqrt{2}$ [111]. As shown in [48], this Bell/EPR game is also an LOCC protocol: Alice and Bob manipulate a quantum resource—the sequence of entangled states—and also communicate classically, both to set up the experiment and to analyze its results.

5.4. QRFs, Contextuality, and Asymptotic Entanglement

If the players in a Bell/EPR game are allowed to communicate their detector settings, they can employ their shared entangled state as a secure communication resource; this is the basis for quantum communication and cryptography protocols (see [38] for an extensive review). As entanglement is, effectively, supraclassical correlation, players of games in which entanglement or other non-local resources can be employed can achieve payoffs larger than those possible in classical versions of those same games. This leads to the concept of a *quantum (or no-signaling) Nash equilibrium* [109] for games that allow use of quantum resources. Note that as discussed in Section 3 above, systems cannot, in general, determine by observation whether they are entangled with their environments [59]; hence, players cannot determine by observation whether the game they are playing is quantum or classical. As the asymptotic state of the quantum FEP is entanglement [29], this immediately implies that each player cannot determine whether they have reached a quantum equilibrium with their environment, or have not.

While in classical uncoupled, or Bayesian, games the concept of "turn taking" can often be elided, in quantum games this is generally not the case. The reason for this is clear in the example of a Bell/EPR game: if Alice uses the basis $(\uparrow, \downarrow)$ while Bob uses $(\nearrow, \searrow)$, different orders of measurement will result in different projections of the shared entangled state. This is an example of operator non-commutativity: measurements along different spin directions do not commute, just as measurements of position and momentum do not commute (both are examples of Heisenberg's uncertainty principle). Operator non-commutativity generically induces quantum contextuality [45,59], defined as the non-causal dependence of one measurement outcome on what others are performed simultaneously [107,112,113]. Contextuality is generally recognized as a resource for both quantum information and complexity [114].

In the generic language employed in Section 2, contextuality can be expressed as non-commutativity between QRFs [45,48,59]. If Q_i and Q_j are non-commuting QRFs with outcome probability distributions P_i and P_j, respectively, when acting on some state ψ, then contextuality manifests as the non-existence of the joint probability distribution $P_i P_j$ (i.e., non-commutativity implies violation of the Kolmogorov axioms; see [115] for a comprehensive review). In the presence of contextuality, joint probabilities over classical strategies can fail to be well defined, hence the shifts in corresponding NEs cannot be detected [116].

6. Discussion

6.1. Rationality

In areas where GT is applied extensively, such as throughout economics, logistics, and the behavioral sciences, it is an overriding assumption that game players act rationally towards optimizing their eventual payoffs given the influence of environmental factors, be these local or global. It is a fundamental principle of GT that when prediction and rationality are aptly combined, an NE is attained. But the question remains, however, in out-of-equilibrium conditions, can rational players successfully predict their opponents' behavior? We can reasonably think, and indeed it is hypothesized in [117], that there is inherent tension between rationality and prediction when players are uncertain of their opponents' *modus operandi* towards payoff. In fact, Foster and Young [117] prove the existence of games in which it is impossible for perfectly rational players to (even approximately) predict the future behavior of their opponents, regardless of any learning rules adopted. There are several slants on this, as investigated in GT and economics (and likely applicable elsewhere too). In [118] (Theorem 1), rational agents in the sense of economics are equivalent to suitably indexed TMs. This implies that decision processes, as implemented by such rational agents, are equivalent to the computing behavior of a suitably indexed TM, and indeed, rational choice, as understood as maximizing choice, is undecidable [118] (Theorem 2). We could follow, e.g., Ewerhart [119] and consider rational players as basing their decisions solely on the provable implications for their assumptions, and the existence of undecidable statements in GT as causing undefinability of rational concepts on the basis of distinctions in logical behavior, e.g., truth versus probability. So, this demands a definition for a rational strategy: a strategy is *rational* if it is equivalent to a best reply to a Bayesian belief, or more generally, if it is a best reply to some lexicographical probability system that satisfies certain consistency conditions, while noting that there are games and strategies for which it is undecidable if they can arise from perfectly rational behavior or from unique predictions, and given the irregularity of beliefs and assumptions that players may have about each other's motives [120].

Assumptions of rationality can in some cases be related to deep assumptions in mathematical logic. There are, for example, statements about two-player, zero-sum games that are undecidable; one such is analogous to the continuum hypothesis (CH) in Zermelo–Fraenkel axiom of choice theory, established in [121]; we refer the reader to [15,122–124] for discussion of this issue from the perspective of Gödel's theorem. As another example, one could take two-person games with pure strategies and for $i = 1, 2$, consider player i with belief set $\Delta_i(\mathfrak{g})$, given a game $\mathfrak{g}$ implementing an infinite-regress logic, denoted EIR^2. Then, for an unsolvable game $\mathfrak{g}$, the theory $(\mathrm{EIR}^2, \Delta_i(\mathfrak{G}))$ is incomplete [125]. As pointed out in [125], we could think of this as a case of self-referentiality, but the actual source of incompleteness is a discrepancy arising from the collective independence of payoffs, predictions, and decision making.

There is also the question of to what extent (Hamiltonian) chaos influences players' behavior towards whether they play rationally or not. In types of simple games (such as rock-paper-scissor), no strategy is seen to be dominant, and in particular, no pure strategy to NEs can exist [126]. The main point is that, in such simple games, it is often the case that questionable psychological heuristics on behalf of the players inevitably suppress any attempt at rational learning to the extent of nonconvergence to an NE. The viewpoint of [126] is simply this: it can be summarized by saying that chaos is a necessary condition for intelligent players to fail attaining to an NE, and the presence of chaos suggests that playing rationally is not always a feasible assumption. Overall, our account reflects upon the prevalence of cognitive bias in many shapes and forms, to the extent that, for the best part, humans are never really close to being (Bayesian) rational game players [127]. On the more technical side, classical economics asserts the existence of at least one NE engaged in strategies, but generally, multiple equilibria are more likely to be the case. Capping this, it can be computationally intractable for players to strictly conform to GT principles in economics [14].

6.2. Entropy of Quantum Games

Since information permeates through this whole circle of ideas, let us comment on how the statistical physics of information/entropy accounts for rational choices of strategies (or the sheer lack of them) in quantum games. The amount of information that a player can obtain about their opponent depends on maximum/minimum entropy criteria, and the rationality of the players in assimilating this information during the course of the game is determined by the prevailing entropy. To see this, let us recall some basic concepts of statistical physics. A (positively valued) *density operator* ρ specifies a mixed ensemble in which each member has an assigned probability of being in a determined state. Its *von Neumann entropy*,

$$S(\rho) = -\mathrm{Tr}\{\rho \ln \rho\} \tag{6}$$

is a probability distribution function [128]. In [129], $S(\rho)$ is maximized subject to $\delta\mathrm{Tr}(\rho) = 0$, and the internal energy constraint $\delta(E) = 0$, leading to

$$\rho_{ii} = \frac{\exp(-\beta E_i)}{\sum_k \exp(-\beta E_k)} \tag{7}$$

where β denotes a thermodynamic parameter (see below). Without the internal energy constraint $\delta(E) = 0$, we have $\rho_{ii} = N^{-1}$, where $N > 0$ is the 'population size'. Effectively $\beta = T^{-1}$, an inverse temperature, and from this there are two cases [129]: (i) $\beta \longrightarrow 0$ is the high temperature limit, in which a canonical ensemble becomes a completely random ensemble; and (ii) $\beta \longrightarrow \infty$ is the low temperature limit, in which a canonical ensemble becomes a pure ensemble where only the ground state is populated. Seeing β as related to the temperature of a statistical system, it can on this account be interpreted as a measure of the *rationality* of the players in question. So, from this statistical physics point of view, we have the expected consequence:

high entropy $\Longleftrightarrow$ *low rationality in the players' behavior*.

6.3. Alternative Equilibria

We have adopted the NE as a focal point of this paper, as the NE has been an idealized central concept of GT that has provided long-standing analytic methods of paramount importance. We acknowledge, however, that particular types of experimental data, as this pertains to "noisy" games with possibly "irrational" players, can escape this analysis, and hence, over the years alternative forms of equilibria have been introduced. One of these is the *quantal response equilibrium* (QRE) of [130], which is more general than the NE in that it relaxes the assumption of best response (to a 'probabilistic' response) and allows noisy optimizing behavior while maintaining consistency of rational expectations. Attaining a QRE entails elements of Bayesian and stochastic choice (e.g., in biological systems and neuroscience) [131,132]. However, the QRE is seen to converge to an NE as the quantal response functions steepen toward approximate best response functions, and in a theoretical framework may not alter predictions as determined by NEs [133]. Further, there are experiments for which QRE solutions do not outperform those of NEs (see, e.g., [134]). Other alternative GT equilibrium theories such as noisy belief and random belief are reviewed in [131].

7. Conclusions

We have shown that the FEP describes generic interactions between physical systems as games in which VFE minimization is the payoff function. Physical interactions are Bayesian games in which powerful no-go theorems restrict what players can know about their own strategies as well as those of their opponents in this context. We have investigated convergence for generic games, and have shown that the classical notions of good regulation, identical synchronization, and winning a GIG are all, in principle, idealizations. We have reviewed quantum games and shown how they implement LOCC protocols.

Undecidability is pervasive in GT; indeed the results reviewed here suggest that all decidable games involve the assumption of knowledge that cannot, as a matter of principle, be obtained by finite observation. Both the undecidability of the frame problem [59,78] and the undecidability of whether two agents are deploying the same QRFs [48] strongly support this conjecture. Arguments to the effect that undecidability is ubiquitous in physics have previously been advanced by Wolfram [135] and Hawking [136].

As all physical systems are, at bottom, quantum systems, generic physical interactions are not just games, but quantum games. To the extent that they involve classical communication—and they must, to be considered games at all—they are instances of LOCC protocols. Their dynamics depend, therefore, on the extent to which quantum resources are manipulated in a coordinated manner by the players. The extent to which such coordination can be either arranged via classical communication or deduced by observation is undecidable [48].

We can conclude, therefore, that GT is much broader in scope than it is often regarded as being: GT's scope includes most, if not all, of physics. Physical systems can, therefore, be regarded as game-playing agents. That this should be the case follows, indeed, from Conway and Kochen's famed "free will" theorem, which shows that no physical system, at any scale, can be fully described by any locally deterministic theory [137].

Author Contributions: Conceptualization, C.F. and J.F.G.; formal analysis, C.F. and J.F.G.; writing—original draft preparation, C.F. and J.F.G.; writing—review and editing, C.F. and J.F.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: All data are contained in the paper.

Acknowledgments: The authors wish to thank two anonymous referees for their comments and suggestions, which were helpful towards the overall presentation of ideas.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CA	Cellular automaton
CCCD	Cone–cocone diagram
CH	Continuum hypothesis
CHSH	Clauser–Horne–Shimony–Holt
EPR	Einstein–Podolsky–Rosen
FEP	Free energy principle
GIG	Generalized imitation game
GT	Game theory
HP	Holographic principle
I/O	Input/output
IPD	Iterated prisoner's dilemma
KL	Kullback–Leibler
LOCC	Local operations and classical communication
MB	Markov blanket
NE	Nash equilibrium
NP	Nondeterministic polynomial
PD	Prisoner's dilemma
PPAD	Polynomial parity arguments on directed graphs
QRE	Quantal response equilibrium
QRF	Quantum reference frame
RNN	Recurrent neural network
SPD	Spatialized prisoner's dilemma
TM	Turing machine
TQFT	Topological quantum field theory
VFE	Variational free energy

References

1. von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior*; Princeton University Press: Princeton, NJ, USA, 1944.
2. Nash, J.F. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* **1950**, *36*, 48–49. [CrossRef] [PubMed]
3. Nash, J.F. Non-cooperative Games. *Ann. Math.* **1951**, *54*, 286–295. [CrossRef]
4. Holt, C.A.; Roth, A.E. The Nash equilibrium: A perspective. *Proc. Natl. Acad. Sci. USA* **2004**, *101*, 3999–4002. [CrossRef] [PubMed]
5. Jacobsen, H.J. On the foundations of Nash equilibrium. *Econ. Phil.* **1996**, *12*, 67–88. [CrossRef]
6. Sethi, R.; Weibull, J. What Is... Nash Equilibrium? *Not. Amer. Math. Soc.* **2016**, *63*, 526–528. [CrossRef]
7. Maynard Smith, J. *Evolution and the Theory of Games*; Cambridge University Press: Cambridge, UK, 1982.
8. Hart, S.; Mas-Colell, A. Uncoupled dynamics do not lead to Nash equilibrium. *Am. Econ. Rev.* **2003**, *93*, 1830–1836. [CrossRef]
9. Milionis, J.; Papdimitriou, C.; Piliouras, G.; Spendlove, K. An impossibility theorem in game dynamics. *Proc. Natl. Acad. Sci. USA* **2023**, *120*, e2305349120. [CrossRef] [PubMed]
10. Sanders, J.B.T.; Farmer, J.D.; Galla, T. The prevalence of chaotic dynamics in games with many players. *Nat. Sci. Rep.* **2018**, *8*, 4902. [CrossRef]
11. Pangallo, M.; Heinrich, T.; Farmer, J.D. Best reply structure and equilibrium convergence in generic games. *Sci. Adv.* **2019**, *5*, eaat1328. [CrossRef]
12. Du, Y. On the complexity of deciding degeneracy in a bimatrix game with sparse payoff matrix. *Theor. Comp. Sci.* **2013**, *472*, 104–109. [CrossRef]
13. Daskalakis, P.G.; Papadimitriou, C. The complexity of computing a Nash equilibrium. *SIAM J. Comp.* **2009**, *39*, 195–259. [CrossRef]
14. Aaronson, S. Why philosophers should care about computational complexity. In *Computability: Turing, Gödel, Church, and Beyond*; Copeland, B.J., Posy, C.J., Shagrir, O., Eds.; MIT Press: Cambridge, MA, USA, 2013; pp. 261–327.
15. Gödel, K. Über formal unentscheidbare sätze der Principia Mathematica und verwandter systeme, I. *Monatsh. Math. Phys.* **1931**, *38*, 173–198. [CrossRef]
16. Tsuji, M.; Da Costa, N.C.A.; Doria, F.A. The incompleteness of theories of games. *J. Philos. Logic* **1998**, *27*, 553–568. [CrossRef]
17. Grim, P. The undecidability of the spatialized prisoner's dilemma. *Theory Decis.* **1997**, *42*, 53–80. [CrossRef]

18. Moore, E.F. Gedankenexperiments on sequential machines. In *Autonoma Studies*; Shannon, C.E., McCarthy, J. Eds.; Princeton University Press: Princeton, NJ, USA, 1956; pp. 129–155.
19. Turing, A.M. Computing machines and intelligence. *Mind* **1950**, *59*, 433–460. [CrossRef]
20. Sato, Y.; Ikegami, T. Undecidability in the imitation game. *Minds Mach.* **2004**, *14*, 133–143. [CrossRef]
21. Milnor, J. *Games Against Nature*; RAND Corp.: Santa Monica, CA, USA, 1951.
22. Feynman, R.P. *Statistical Mechanics*; Benjamin: Reading, MA, USA, 1972.
23. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138. [CrossRef]
24. Friston, K. Life as we know it. *J. R. Soc. Interface* **2013**, *10*, 20130475. [CrossRef]
25. Friston, K.J. A free energy principle for a particular physics. *arXiv* **2019**, arXiv:1906.10184.
26. Ramstead, M.J.; Sakthivadivel, D.A.R.; Heins, C.; Koudahl, M.; Millidge, B.; Da Costa, L.; Klein, B.; Friston, K.J. On Bayesian mechanics: A physics of and by beliefs. *R. Soc. Interface Focus* **2023**, *13*, 20220029. [CrossRef]
27. Friston, K.J.; Da Costa, L.; Sakthivadivel, D.A.R.; Heins, C.; Pavliotis, G.A.; Ramstead, M.J.; Parr, T. Path integrals, particular kinds, and strange things. *Phys. Life Rev.* **2023**, *47*, 35–62. [CrossRef] [PubMed]
28. Horsman, C.; Stepney, S.; Wagner, R.C.; Kendon, V. When does a physical system compute? *Proc. R. Soc. A* **2014**, *470*, 20140182. [CrossRef] [PubMed]
29. Fields, C.; Friston, K.J.; Glazebrook J.F.; Levin M. A free energy principle for generic quantum systems. *Prog. Biophys. Mol. Biol.* **2022**, *173*, 36–59. [CrossRef]
30. Fields, C.; Fabrocini, Friston, K.; Glazebrook, J.F.; Hazan, H.; Levin, L.; Marcianò, A. Control flow in active inference systems, Part I: Formulations of classical and quantum active inference. *IEEE Trans. Mol. Biol. Multi-Scale Commun.* **2023**, *9*, 235–245. [CrossRef]
31. Fields, C.; Glazebrook, J.F. Representing measurement as a thermodynamic symmetry breaking. *Symmetry* **2020**, *12*, 810. [CrossRef]
32. Addazi, A.; Chen, Fabrocini, F.; Fields, C.; Greco, E.; Lulli, M.; Marcianò, A.; Pasechnik, R. Generalized holographic principle, gauge invariance and the emergence of gravity à la Wilczek. *Front. Astron. Space Sci.* **2021**, *8*, 563450. [CrossRef]
33. Fields, C.; Glazebrook, J.F.; Marcianò, A. Reference frame induced symmetry breaking on holographic screens. *Symmetry* **2021**, *13*, 408. [CrossRef]
34. 't Hooft, G. Dimensional reduction in quantum gravity. In *Salamfestschrift*; Ali, A., Ellis, J., Randjbar-Daemi, S., Eds.; World Scientific: Singapore, 1993; pp. 284–296.
35. Susskind, L. The world as a hologram. *J. Math. Phys.* **1995**, *36*, 6377–6396. [CrossRef]
36. Bousso, R. The holographic principle. *Rev. Mod. Phys.* **2002**, *74*, 825–874. [CrossRef]
37. Fields, C.; Glazebrook, J.F.; Marcianò, A. The physical meaning of the holographic principle. *Quanta* **2022**, *11*, 72–96. [CrossRef]
38. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*; Cambridge University Press: New York, NY, USA, 2000.
39. Wheeler, J.A. Information, physics, quantum: The search for links. In *Complexity, Entropy, and the Physics of Information*; Zurek, W., Ed.; CRC Press: Boca Raton, FL, USA, 1989; pp. 3–28.
40. Ekert, A.K.; Huttner, B.; Palmer, G.M.; Peres, A. Eavesdropping on quantum cryptographical systems. *Phys. Rev. A* **1994**, *50*, 1047–1056. [CrossRef]
41. Pegg, D.; Barnett, S.; Jeffers, J. Quantum theory of preparation and measurement. *J. Mod. Opt.* **2010**, *49*, 913–924. [CrossRef]
42. Aharonov, Y.; Kaufherr, T. Quantum frames of reference. *Phys. Rev. D* **1984**, *30*, 368–385. [CrossRef]
43. Bartlett, S.D.; Rudolph, T.; Spekkens, R.W. Reference frames, superselection rules, and quantum information. *Rev. Mod. Phys.* **2007**, *79*, 555–609. [CrossRef]
44. Fields, C.; Glazebrook, J.F. A mosaic of Chu spaces and Channel Theory I: Category-theoretic concepts and tools. *J. Expt. Theor. Artif. Intell.* **2019**, *31*, 177–213. [CrossRef]
45. Fields, C.; Glazebrook, J.F. Information flow in context-dependent hierarchical Bayesian inference. *J. Expt. Theor. Artif. intell.* **2022**, *34*, 111–142. [CrossRef]
46. Barwise, J.; Seligman, J. *Information Flow: The Logic of Distributed Systems*; Cambridge University Press, Cambridge, UK, 1997.
47. Fields C.; Glazebrook, J.F.; Marcianò, A. Sequential measurements, topological quantum field theories, and topological quantum neural networks. *Fortschr. Phys.* **2022** *70*, 2200104. [CrossRef]
48. Fields, C.; Glazebrook, J.F.; Marcianò, A. Communication protocols and QECCs from the perspective of TQFT, Part I: Constructing LOCC protocols and QECCs from TQFTs. *Fortschr. Phys.* **2024**, *72*, 202400049. [CrossRef]
49. Pearl, J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*; Morgan Kaufmann: San Mateo, CA, USA, 1988.
50. Hutter, M. *Universal Artificial Intellegence Sequential Decisions Based on Algorithmic Probability*; Springer: Berlin/Heidelberg, Germany, 2006.
51. MacKay, D.J. Free-energy minimisation algorithm for decoding and cryptoanalysis. *Electron. Lett.* **1995**, *31*, 445–447. [CrossRef]
52. Ruffini, G. An algorithmic information theory of consciousness. *Neurosci. Cons.* **2017**, *2017*, nix019. [CrossRef]
53. Wallace, C.S.; Dowe, D.L. Minimum message length and Kolmogorov complexity. *Comput. J.* **1999**, *42*, 270–283. [CrossRef]
54. Hoffman, D.D.; Singh, M.; Prakash, C. The interface theory of perception. *Psychon. Bull. Rev.* **2015**, *22*, 1480–1506. [CrossRef] [PubMed]
55. Rumelhart, D.E.; Hinton, G.E.; Williams, R.J. Learning representations by back-propagating errors. *Nature* **1986**, *323*, 533–536. [CrossRef]
56. Oudeyer, P.-Y.; Kaplan, F.; Hafner, V. Intrinsic motivation systems for autonomous mental development. *IEEE Trans. Evol. Comput.* **2007**, *11*, 265–286. [CrossRef]

57. Gopnik, A. Explanation as orgasm and the drive for causal understanding: The evolution, function and phenomenology of the theory-formation system. In *Cognition and Explanation*; Keil, F., Wilson, R., Eds.; MIT Press: Cambridge, MA, USA, 2000; pp. 299–323.
58. Fields, C.; Glazebrook, J.F.; Levin, M. Principled limitations on self-representation for generic physical systems. *Entropy* **2024**, *26*, 194. [CrossRef]
59. Fields, C.; Glazebrook, J.F. Separability, contextuality, and the quantum Frame Problem. *Int. J. Theor. Phys.* **2023**, *62*, 159. [CrossRef]
60. McCarthy, J.; Hayes, P. J. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*; Michie, D., Meltzer, B., Eds.; Edinburgh University Press: Edinburgh, UK, 1969; Volume 4, pp. 463–502.
61. Harsanyi, J. C. Games with incomplete information played by Bayesian players, Part I. *Manag. Sci.* **1967**, *14*, 159–183. [CrossRef]
62. Quine, W.V.O. *Word and Object*; MIT Press: Cambridge, MA, USA, 1960.
63. Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; Pezzulo, G. Active inference: A process theory. *Neural Comput.* **2017**, *29*, 1–49. [CrossRef] [PubMed]
64. Ramstead, M.J.D.; Badcock, P.B.; Friston, K.J. Answering Schrödinger's question: A free-energy formulation. *Phys. Life Rev.* **2018**, *24*, 1–16. [CrossRef]
65. Ramstead, M.J.D.; Constant, A.; Badcock, P.B.; Friston, K.J. Variational ecology and the physics of sentient systems. *Phys. Life Rev.* **2019**, *31*, 188–205. [CrossRef]
66. Raup, D.M. *Extinction: Bad Genes or Bad Luck?*; Norton: New York, NY, USA, 1991.
67. Friston, K.J.; Frith, C.D. Duet for one. *Conscious. Cogn.* **2015**, *36*, 390–405. [CrossRef] [PubMed]
68. Friston, K.J.; Frith, C.D. Active inference, communication and hermeneutics. *Cortex* **2015**, *68*, 129–143. [CrossRef]
69. Schiff, S.J.; So, P.; Chang, T.; Burke, R.E.; Sauer, T. Detecting dynamical interdependence and generalized synchrony through mutual prediction in a neural ensemble. *Phys. Rev. E* **1996**, *54*, 6708–6724. [CrossRef] [PubMed]
70. Pecora, L.M.; Carroll, T.L.; Johnson, G.A.; Mar, D.J.; Heagy, J.F. Fundamentals of synchronization in chaotic systems, concepts, and applications. *Chaos* **1997**, *7*, 520–543. [CrossRef]
71. Friston, K.; Breakspear, M.; Deco, G. Perception and self-organized instability. *Front. Comp. Neurosci.* **2012**, *6*, 44. [CrossRef]
72. Friston, K.; Sengupta, B.; Auletta, G. Cognitive dynamics: From attractors to active inference. *Proc. IEEE* **2014**, *102*, 427–445. [CrossRef]
73. Palacios, E.R.; Isomura, T.; Parr, T.; Friston, K. The emergence of synchrony in networks of mutually inferring neurons. *Nature Sci. Rep.* **2019**, *9*, 6412. [CrossRef] [PubMed]
74. Bilek, E.; Zeidman, P.; Kirsch, P.; Tost, H.; Meyer-Lindenberg, A.; Friston, K. Directed coupling in multi-brain networks underlies generalized synchrony during social exchange. *NeuroImage* **2022**, *252*, 119038. [CrossRef]
75. Rice, H.G. Classes of recursively enumerable sets and their decision problems. *Trans. Am. Math. Soc.* **1953**, *74*, 358–366. [CrossRef]
76. Turing, A. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.* **1937**, *42*, 230–265. [CrossRef]
77. Hopcroft, J.E.; Ullman, J.D. *Introduction to Automata Theory, Languages, and Computation*; Addison-Wesley: Boston, MA, USA, 1979.
78. Dietrich, E.; Fields, C. Equivalence of the Frame and Halting problems. *Algorithms* **2020**, *13*, 175. [CrossRef]
79. Dunbar, R.I.M. The social brain: Mind, language and society in evolutionary perspective. *Annu. Rev. Anthropol.* **2003** *32*, 163–181. [CrossRef]
80. Eisert, J.; Wilkens, M.; Lewenstein, M. Quantum games and quantum strategies. *Phys. Rev. Lett.* **1999**, *83*, 3077–3080. [CrossRef]
81. Myerson, R.B. *Game Theory: An Analysis of Conflict*; MIT Press: Cambridge, MA, USA, 1991.
82. Chater, N. *The Mind is Flat*; Allen Lane: London, UK, 2018.
83. Markose, S.M. Complex type 4 structure changing dynamics of digital agents: Nash equilibria of a game with arms race innovations. *J. Dyn. Games Am. Inst. Math. Sci.* **2017**, *4*, 255–284. [CrossRef]
84. Prokopenko, M.; Harré, M.; Lizier, J.T.; Boschetti, F.; Pappas, P.; Kauffmann, S. Self-referential basis of undecidable dynamics: from the Liar Paradox and the Halting Problem to the Edge of Chaos. *Phys. Life Rev.* **2019**, *31*, 134–156. [CrossRef]
85. Taiji, M.; Ikegami, T. Dynamics of internal models in game players. *Phys. D Nonlinear Phenom.* **1999**, *134*, 253–266. [CrossRef]
86. Ikegami, T.; Taiji, M. *Uncertainty, Possible Worlds and Coupled Dynamical Recognizers*; University of Tokyo: Tokyo, Japan, 1998.
87. Eisert, J.; Wilkens, M. Quantum games. *J. Modern Opt.* **2000**, *47*, 2543–2566. [CrossRef]
88. Neyman, A. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Econ. Lett.* **1985**, *19*, 227–229. [CrossRef]
89. Conant R.C.; Ashby, W.R. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* **1970**, *1*, 89–97. [CrossRef]
90. Worden, R.P. The requirement for cognition, in an equation. *arXiv* **2024**, arXiv:2405.08601.
91. Kari, J. Rice's theorem for the limit sets of cellular automata. *Theor. Comp. Sci.* **1994**, *127*, 229–254. [CrossRef]
92. Di Lena, P.; Margara, L. On the undecidability of the limit behavior of Cellular Automata. *Theor. Comp. Sci.* **2020**, *411*, 1075–1084. [CrossRef]
93. Meyer, D.A. Quantum strategies. *Phys. Rev. Lett.* **1999**, *82*, 10520–1055. [CrossRef]
94. Vos Fellman, P.; Vos Post, J. Quantum Nash equiibria and quantum computing. In *Unifying Themes in Complex Systems*; Minai, A., Braha, D., Bar-Yam, Y., Eds.; Springer: Berlin/Heidelberg, Germany, 2010; pp. 454–461.
95. Scarani, V.; Iblisdir, S.; Gisin, N.; Acin, A. Quantum cloning. *Rev. Mod. Pys.* **2005**, *77*, 1225–1256. [CrossRef]
96. Ney, P.-M.; Notarnicola, S.; Montanegro, S.; Morigi, S. Entanglement in the quantum Game of Life. *Phys. Rev. A* **2022**, *105*, 012416. [CrossRef]
97. Kaur, H.; Kumar, A. Analysing the role of entanglement in the three-qubit Vaidman's game. In Proceedings of the 2017 International Conference on Intelligent Communication and Computational Techniques (ICCT), Jaipur, India, 22–23 December 2017; pp. 96–101.
98. Chitambar, E.; Leung, D.; Mančinska, L.; Ozols, M.; Winter, A. Everything you always wanted to know about LOCC (but were afraid to ask). *Comms. Math. Phys.* **2014**, *328*, 303–326. [CrossRef]

99. Zurek, W.H. Environment-induced superselection rules. *Phys. Rev. D* **1982**, *26*, 1862–1880. [CrossRef]
100. Joos, E.; Zeh, H.D. The emergence of classical properties Through interaction with the environment. *Zeitschr. Phys. B* **1985**, *59*, 233–243. [CrossRef]
101. Zurek, W.H. Decoherence, einselection, and the quantum origins of the classical. *Rev. Mod. Phys.* **2003**, 715–775. [CrossRef]
102. Blume-Kohout, R.; Zurek, W.H. Quantum Darwinism: Entanglement, branches, and the emergent classicality of redundantly stored quantum information. *Phys. Rev. A* **2006**, *73*, 062310. [CrossRef]
103. Zurek, W.H. Quantum Darwinism. *Nature Phys.* **2009**, *5*, 181–188. [CrossRef]
104. Aspect, A.; Grangier, P.; Roger, G. Experimental tests of realistic local theories via Bell's theorem. *Phys. Rev. Lett.* **1981**, *47*, 460–463. [CrossRef]
105. Georgescu, I. How the Bell tests changed quantum physics. *Nat. Phys.* **2021**, *3*, 374–376. [CrossRef]
106. Bell, J.S. On the Einstein-Podolsky-Rosen paradox. *Physics* **1964**, *1*, 195–200. [CrossRef]
107. Bell, J.S. On the problem of hidden variables in quantum mechanics. *Rev. Mod. Phys.* **1966**, *38*, 447–452. [CrossRef]
108. Mermin, N.D. Is the Moon there when nobody looks? Reality and the quantum theory. *Phys. Today* **1985**, *38*, 38–47. [CrossRef]
109. Brunner, N.; Linden, N. Connection between Bell nonlocality and Bayesian game theory. *Nat. Commun.* **2013**, *4*, 2057. [CrossRef]
110. Clauser, J.F.; Horne, M.A.; Shimony, A.; Holt, R.A. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **1969**, *23*, 880–884. [CrossRef]
111. Cirelson, B.S. Quantum generalizations of Bell's inequality. *Lett. Math. Phys.* **1980**, *4*, 93–100. [CrossRef]
112. Kochen, S.; Specker, E.P. The problem of hidden variables in quantum mechanics. *J. Math. Mech.* **1967**, *17*, 59–87. [CrossRef]
113. Mermin, N.D. Hidden variables and the two theorems of John Bell. *Rev. Mod. Phys.* **1993**, *65*, 803–815. [CrossRef]
114. Howard, M.; Wallman, J.; Veitch, V.; Emerson, J. Contextuality supplies the 'magic' for quantum computation. *Nature* **2014**, *510*, 351–686. [CrossRef] [PubMed]
115. Khrennikov, A. Contextuality, complementarity, signaling, and Bell tests. *Entropy* **2022**, 24, 1380. [CrossRef] [PubMed]
116. Fourny, G. On the interpretation of quantum theory as games between physicists and nature played in Minkowski spacetime. *arXiv* **2024**, arXiv:2405.20143.
117. Foster, D.P.; Young, H.P. On the impossibility of prediction of the behavior of rational agents. *Proc. Natl. Acad. Sci. USA* **2001**, *98*, 12848–12853. [CrossRef] [PubMed]
118. Velupillai, K.V. Uncomputability and undecidability in economic theory. *Appl. Math. Comput.* **2009**, *215*, 1404–1416. [CrossRef]
119. Ewerhart, C. On Strategic Reasoning and Theories of Rational Behavior. Ph.D. Thesis, University of Bonn, Bonn, Germany, 1997.
120. Ewerhart, C. Rationality and the definition of consistent pairs. *Int. J. Game Theory* **1998**, *27*, 49–59. [CrossRef]
121. Fey, M. An undecidable statement regarding zero-sum games. *Games Econ. Behav.* **2024**, *145*, 19–26. [CrossRef]
122. Gödel, K. *The Consistency of the Axiom of Choice and of the Generalized Continuum-Hypothesis with the Axioms of Set Theory*; Princeton University Press: Princeton, NJ, USA, 1940; Volume 3; Annals of Mathematics Studies.
123. Cohen, P.J. The independence of the Continuum Hypothesis, Part I. *Proc. Nat. Acad. Sci. USA* **1963**, *50*, 1143–1148. [CrossRef]
124. Cohen, P.J. The independence of the Continuum Hypothesis, Part II. *Proc. Nat. Acad. Sci. USA* **1964**, *51*, 105–110. [CrossRef]
125. Hu, T.-W.; Kaneko, M. *Game Theoretic Decidability and Undecidability*; WINPEC Working Paper Series No. E1410; Social Science Research Network (SSRN): Rochester, NY, USA, 2015.
126. Sato, Y.; Akiyama, E.; Farmer, J.D. Chaos in learning a simple two-person game. *Proc. Natl. Acad. Sci. USA* **2002**, *99*, 4748–451. [CrossRef]
127. Kahneman, D.; Slovic, P.; Tversky, A. *Judgement under Uncertainty: Heuristics and Biases*; Cambridge University Press: Cambridge, UK, 1982.
128. von Neumann, J. Thermodynamik quantummechanischer Geshameiten. *Gött. Nach.* **1927**, *1*, 273–291.
129. Hidalgo, E.G. Quantum games entropy. *Phys. A Stat. Mech. Appl.* **2017**, *383*, 797–804. [CrossRef]
130. McKelvey, R.D.; Palfrey, R.D. Quantal response equilbria for normal form games. *Games Econ. Behav.* **1995**, *10*, 6–38. [CrossRef]
131. Friedman, E. Stochastic equilibria: Noise in actions or beliefs? *Am. Econ. J. Microeconomics* **2022**, *14*, 94–142. [CrossRef]
132. Bland, J.R. Bayesian inference for Quantal Response Equilibrium in normal-form games. *Games Econ. Behav.* 2024, *in press*.
133. Goeree, J.K.; Holt, C.A.; Palfrey, T.R. Quantal response equilibria. In *Behavioural and Experimental Economics*; Durlauf, S.N., Blume, L.E., Eds.; The New Palgrave Economics Collection; Palgrave Macmillan London: London, UK, 2010; pp. 234–242.
134. Volacu, A. Mixed strategy Nash Equilibrium and Quantal Response Equilibrium: An experimental comparison using RPS games. *Theor. Appl. Econ.* **2014**, *XX1*, 89–118.
135. Wolfram, S. Undecidability and intractability in theoretical physics. *Phys. Rev. Lett.* **1985**, *54*, 735–738. [CrossRef] [PubMed]
136. Hawking, S. Gödel and the End of Physics. Lecture at the Dirac Centennial Celebration. Centre for Mathematical Sciences, University of Cambridge, Cambridge, UK, 2002. Available online: https://www.damtp.cam.ac.uk/events/strings02/dirac/hawking.html (accessed on 13 January 2024).
137. Conway, J.H.; Kochen, S. The strong free will theorem. *Not. AMS* **2009**, *56*, 226–232.