

Review

Information Theory for Agents in Artificial Intelligence, Psychology, and Economics

Michael S. Harré¹  0000-0003-2199-2515¹ Affiliation: michael.harre@sydney.edu.au

† Current address: Complex Systems Research Group, Faculty of Engineering, The University of Sydney, Sydney, Australia, 2006

Abstract: This review covers some of the central relationships between artificial intelligence, psychology, and economics through the lens of information theory, specifically focusing on formal models of decision-theory. In doing so we look at a particular approach that each field has adopted and how information theory has informed the development of the ideas of each field. A key theme is expected utility theory, its connection to information theory, and the Bayesian approach to decision-making and forms of (bounded) rationality. What emerges from this review is a broadly unified formal perspective derived from three very different starting points that reflect the unique principles of each field. Each of the three approaches reviewed can, in principle at least, be implemented in a computational model in such a way that, with sufficient computational power, they could be compared with human abilities in complex tasks. However, a central critique that can be applied to all three approaches was first put forward by Savage in *The Foundations of Statistics* and recently brought to the fore by the economist Binmore: Bayesian approaches to decision-making work in what Savage called ‘small worlds’ but cannot work in ‘large worlds’. This point, in various different guises, is central to some of the current debates about the power of artificial intelligence and its relationship to human-like learning and decision-making. Recent work on artificial intelligence has gone some way to bridging this gap but significant questions still need to be answered in all three fields in order to make progress on these problems.

Keywords: Decision Theory, Psychology, Artificial Intelligence, Economics, Artificial Neural Networks, Neuroscience, Information Theory, Bayesian Learning, Free Energy

1. Introduction

The goal of this article is to review a relatively unexplored but interesting intersection between three different research areas that motivated the special issue this article appears in: economics, artificial intelligence, and psychology. One way in which these fields intersect with one another is via decision theory, each bringing something different to the discussion but with points of tension in how they conceptualise the issues. In highlighting a subset of the differences and similarities, this article aims to provide a perspective that is not well covered in the literature. To do this we review these fields from an information theoretical perspective, chosen because of its popularity as an empirical tool, but also because it has noteworthy properties in the context of decision theory, providing a overarching technical framework. Each field also has a relatively well understood path to the other two, but seen as a whole there is a tension between all three that suggests further exploration, Figure 1 illustrates this unification-tension in decision theory. In this Introduction we very briefly cover some earlier uses of information theory in each field, more details are also in the following sections, before describing how the article is structured. It is hoped that those interested in recent progress in these three fields, how they overlap with one another, their recent progress, and the theoretical concerns in implementing them in simulations will find this article of interest.




Figure 1. A schematic of the relationships between the three fields. Each element has a direct link of the same colour to the other two, but no single path navigates the whole circuit.

One of the most important contributors to the debate regarding entropy and its use in economics comes from the work Georgescu-Roegen [1]. In his extensive body of work Georgescu-Roegen was concerned with the analysis of an entire economic process and the relationship it has with entropy and thermodynamics. At the smaller scale of markets, in the economics of population dynamics, Wilson's entropic methods [2] have been used to study long term population flows using methods based on maximum entropy [3] and Fisher information [4]. The stability of financial markets have been studied using mutual information [5,6] and transfer entropy [7], drawing an analogy from the instability of aggregate economic dynamics and the physics of phase-transitions [8] in which mutual information [9], transfer entropy [10], and Fisher information [11] have been used to examine strongly non-linear systems. In market dynamics Aoki [12] has used maximum entropy methods to model the approach to stationarity including interactions between agents and these same maximum entropy techniques have been used to studied game theory [13,14]. However, despite these diverse applications and both the necessary software [15] and texts introducing the methods [16] (particularly Chapter 6 on finance and economics) there has only been limited take-up of these methods in mainstream economics.

In psychology, the history of information theory is also quite uneven and Laming has written an extensive history and critique [17] of the use of these methods in psychology. Early on the application of Shannon's ideas to psychology were enthusiastic, shortly after the development of information theory in 1948 [18] Miller published his 1956 article *The magical number seven, plus or minus two: Some limits on our capacity for processing information* [19]. In this well-cited article information theory is used to measure *the span of absolute judgement*, what might be called our short-term or working memory today, estimating it to contain around seven items. At about the same time (1954) McGill published his article *Multivariate information transmission* [20] using sophisticated formal methods to analyse multivariate data, stating:

It is not widely understood, however, that the tools made available by communication theory are useful for analyzing data, whether or not we believe that the human organism is best described as a communications system.

So these two papers lay the foundations for psychology to use information theory as both a method for data analysis and as a method for understanding human cognition as information storage or processing. However, as the mathematical psychologist Duncan Luce has written [21], information theory has never played a central role in psychology. Luce's article is critical of the information theory - psychology partnership, in writing on

the psychological studies of mean response times to stimuli, a common experimental paradigm still used today, Luce writes [21, pg. 185]:

In early experiments, mean response time appeared to grow linearly with uncertainty, but glitches soon became evident. The deepest push in the response time direction was Donald Laming's (1968) subtle *Information Theory of Choice-Reaction Times*, although he later stated: "This idea does not work. While my own data (1968) might suggest otherwise, there are further unpublished results that show it to be hopeless" (Laming, 2001, p. 642).

The enthusiasm – nay, faddishness – of the times is hard to capture now. Many felt that a very deep truth of the mind had been uncovered. Yet, Shannon was skeptical [...] Myron Tribus (1979, p. 1) wrote: "In 1961 Professor Shannon, in a private conversation, made it quite clear to me that he considered applications of his work to problems outside of communication theory to be suspect and he did not attach fundamental significance to them." These skeptical views strike me as appropriate.

Luce sees the key issue to be in the hidden structure of human responses to stimuli, but information theory, he argues, is agnostic to these subtleties of cognitively mediated relationships between stimuli. For example he describes the highly sensitive nature of the correlation structure of human response times when there is either short or long separations between the stimuli, implying that they are not separable in a way that information theory requires them to be [21, pg. 185]:

But in my opinion, the most important answer lies in the following incompatibility between psychology and information theory. The elements of choice in information theory are absolutely neutral and lack any internal structure; the probabilities are on a pure, unstructured set whose elements are functionally interchangeable.

However some recent work in the expertise literature has been able to show that, at least at an aggregate data-analysis level, some of the results obtained in the psychology of decision theory have been re-cast in an information theoretical framework [22–24], matching the earlier results that used more conventional methods. In these studies some of the high level qualitative features observed in human behavioural studies can be seen using information theory, suggesting that the data-analytical use of information theory has applications in studies of behavioural psychology. In what follows we will see that, at least to some extent, these concerns are not as insurmountable as Luce or Laming suggest, but that we are also still a long way away from being able to draw definitive conclusions.

Within the field of neuroscience though there has been much greater take-up of information theoretical methods, of which only two are mentioned here. In network inference and data analysis the development of software such as JIDT [15] and IDTxL [25] that can be applied to neuro-physiological data [26–28] has made the application of computationally intensive, information theoretical tools much easier in recent years. In another specialisation, theoretical work on decision theory and cognition [29–32] has made significant inroads in modelling some of the theoretical principles that may underpin decision-making, and the neuro-scientific principles of learning, memory, and behaviour that connect these theories to practical concerns have begun to emerge [33,34].

Finally, in the field of artificial intelligence (AI) methods based on information theory such as mutual information and relative (cross) entropy have had a long and productive history [35]. But a distinction should be made between different schools of thought regarding the purpose of studying artificial intelligence. For example as a model

of human cognition there are a number of quite different mathematical approaches at different levels of analysis and development such as the free energy principle [36] (which is the focus of Section 5), integrated information theory [37], or the AIXI agent for universal intelligence [38]. Each of these approaches attempts to formally describe human intelligence and is less concerned with, although not totally disinterested in, the consequences for AI applied to data analysis.

On the other hand there is an extensive body of work that uses AI as a tool for the analysis of large and complex data-sets and these are less concerned with what this tells us about human cognition and decision-making. However these methods often make use of ideas from information theory in order to improve performance, for example in feature selection tasks mutual information is used to measure the relevance of different features [39] as well as feature selection using entropy-based filters [40], information gain methods [41], one-shot learning [42], and deep learning [43]. In these applications information theory plays an important role, but the emphasis is towards the practical implementation of effective algorithms for data analysis, rather than the foundations of decision-making.

There is a third position between these two ends of the scale that is also worth mentioning because these AIs are opening up the debate on the potential for AI to perform at human like levels on tasks that are complex enough to challenge humans in ways that are quite surprising, albeit still in relatively narrow domains. A candidate for the earliest of the most recent evolution of “big data” AI to challenge human-like reasoning is IBM’s Watson [44], an AI that in 2012 was able to win against past human champions of the US game show *Jeopardy!*, a game where the contestants need to guess the question given the answer, a type of backward induction that requires sophisticated aspects of inference and contextualisation that humans are particularly good at and AIs were particularly poor at. The next candidate is AlphaGo [45] that beat the 9 dan Go champion Lee Sedol in 2016 and its decedent, AlphaZero [46], that learned to play Go, Chess, and Shogi entirely from self-play. The last example is the auto-regressive language model called GPT-3 [47], probably the most controversial of these three AIs because of its ability to generate natural language that can be mistaken for human creative writing and reasoning. For a brief introduction see the OpenAI API website¹ and for a review of GPT-3’s ability to ‘think’ by forming analogies see [48]. These AIs have generated considerable debate as researchers and practitioners work through what is hyperbole and what is reasonable extrapolation of the near future of AI and its ability to perform human-like reasoning, for a recent review of what currently seems plausible in AI and its shortcomings see [49].

In the following Sections we introduce some of the historical background for each field and then focus on a very specific line of research or reasoning in order to illustrate the broader perspective of this article. The reason is a matter of practicality: each field could not be summarised very easily in a single article, although where I know of such an article I have referenced it, and to attempt a full review for all three fields is probably beyond any reasonable length article. So I have selected work that has particularly motivated this special issue, and so it is naturally idiosyncratic for that reason, but there are some principles I have roughly followed to further narrow the field. The work needs to be well-established with substantial prior work to draw on and ample opportunity for it to have been critiqued. It also needs to be expressible formally so that the mathematical connections, if they exist, can be discussed. The final section brings some of the similarities and differences into play in order to illustrate where the lines of tension lie and what this may mean for future work. There are many other ways to have sliced and diced these topics, but I hope the mix that appears here provides an interesting and

¹ <https://openai.com/blog/openai-api/>

illuminating perspective.

2. Economic rationality and expected utility

What we mean by “rationality” and what it means for how we make decisions has a long and extensively debated past, see for example the many different definitions that can be found in [50]. The historical background to these debates is important for understanding the different interpretations of rationality, but for this section economic rationality will mean *rational decision theory* using *expected utility theory* that is summarised in [51] and follows Binmore’s approach [52] that is based on Savage’s original work on Bayesian decision theory [53]. In this section I introduce the axioms of economic decision theory in order to make clear exactly what utility means in economics and how it is interpreted in the mainstream (classical) sense. The key point is that this interpretation of rationality avoids the need to have any representation of the psychological processes that result in a decision being made.

An agent is given a set of lotteries \mathcal{L} for which each $l_i \in \mathcal{L}$ is a random process with j possible outcomes $o_{i,j}$ that occur with probability $p_{i,j}$. We can also combine lotteries in the following way. Say the agent is given two lotteries to choose between, l_1 with probability p and l_2 with probability $1 - p$, then we can define a third lottery as the combination of these two lotteries: $\{pl_1, (1 - p)l_2\}$. We also use the symbol \preceq to represent a weak preference relationship between two lotteries l_1 and l_2 from a set \mathcal{L} . We interpret $l_1 \preceq l_2$ as meaning that l_2 is at least as preferable as l_1 , and the agent can then weakly order the lotteries in \mathcal{L} . With these definitions we can state the necessary axioms regarding lotteries and preferences that are used in establishing Savage’s key theorems :

Axiom 1 Completeness:

For any $l_1, l_2 \in \mathcal{L}$ either $l_1 \preceq l_2$ or $l_2 \preceq l_1$.

Axiom 2 Transitivity:

For any $l_1, l_2, l_3 \in \mathcal{L}$, if $l_1 \preceq l_2$ and $l_2 \preceq l_3$, then $l_1 \preceq l_3$.

In the case where $l_1 \preceq l_2$ and $l_2 \preceq l_1$ an agent is indifferent to these two lotteries and we use the notation: $l_1 \sim l_2$. The last two axioms are:

Axiom 3 Continuity:

Suppose we have $l_1, l_2 \in \mathcal{L}$, then for any $p \in [0, 1]$ there is a third lottery l_3 such that: $\{pl_1, (1 - p)l_2\} \sim l_3$.

Axiom 4 Independence:

Suppose we have $l_1, l_2 \in \mathcal{L}$, then for any l_3 and $p \in [0, 1]$: $\{pl_1, (1 - p)l_3\} \preceq \{pl_2, (1 - p)l_3\}$.

These axioms are sufficient to state and prove one of the key theorems of economic decision theory:

Theorem: Savage [53, Thm. 1 pg. 73 and Thm. 4. pg. 75]

If an agent is asked to choose between two lotteries $l_i, l_j \in \mathcal{L}$ and their behaviour satisfies axioms 1-4 above then there exists real-valued utilities $U(l_i)$ such that:

$$l_i \preceq l_j \iff U(l_i) \leq U(l_j) \quad (1)$$

That is to say that a weakly ordered preference over choices that has been observed in the behaviour of an agent, and without reference to the inner workings of the agent’s

mind, is sufficient to infer a numerical utility function. Finally, choosing a lottery is *rational*, or the agent is said to act rationally if the agent chooses the lottery l_* that would be consistent with maximising their expected utility:

$$l_* = \max_i U(l_i) \quad \forall l_i \in \mathcal{L} \quad (2)$$

This axiomatic approach has played a fundamental role in expected utility theory, but equally important is the subsequent relationship between *utility* and *rationality* that informs the interpretation of economic rationality in the context of psychological or cognitive processes.

Binmore has written that economic theory does not make any claims about the *cognitive processes* that maximise a utility function, and goes further to say that within economics a utility function does not need to exist at all. The following passage is taken from his book *Rational Choice*:

Economists after Bentham [1748 - 1832] became increasingly uncomfortable, not only with the naïve hypothesis that our brains are machines for generating utility, but with all attempts to base economics on psychological foundations. The theory of revealed preference therefore makes a virtue of assuming nothing whatever about the psychological causes of our choice behavior [...] Modern decision theory succeeds in accommodating the infinite variety of the human race within a single theory simply by denying itself the luxury of speculating about what is going on inside someone's head. Instead, it pays attention only to what people do. [52, Chapter 1: Revealed Preferences]

This argument is based on the idea that agents behave *as if* they were being driven by a psychological process that is using an expected utility function. The *as if* argument was introduced in a highly influential article by Milton Friedman on "Positive Economics" [54] and by now the argument has a long history, see for example the recent collection of articles in [55]. Binmore articulated the consequences of *as if* reasoning for revealed preferences in *Rational Choice*:

In revealed-preference theory, it isn't true that [an agent] chooses b rather than a because the utility of b exceeds the utility of a . This is the Causal Utility Fallacy. It isn't even true that [an agent] chooses b rather than a because she prefers b to a . On the contrary, it is because [an agent] chooses b rather than a that we say that [they] prefer b to a , and assign b a larger utility. [52, Chapter 1: Revealed Preferences]

To summarise this view:

The rationality of expected utility theory requires an agent's preferences to be described by axioms 1-4, that they are stable over time, and that they are revealed through the agent's behaviour. From these premises an observer can infer, by Savage's theorems, that an agent's preferences can be represented using a numerical utility function in the form of Equation 1 and that the agent's behaviour is optimal with respect to this inferred utility function. However this does not mean that a utility function exists or is used by the agent to optimise their decision, optimisation is inferred by consistency of preferences, and the agent only behaves as though a utility exists and as though they are using it to optimise the choices they make.

The history of the wide acceptance of expected utility theory in economics and the role of Savage, Friedman, and Samuelson is an interesting story of how the theoretical difficulties that Samuelson saw in expected utility theory were resolved by introducing Axiom 4 (Independence), finally winning Samuelson over to the theory [56]. From

expected utility theory Savage showed how an economic agent can use Bayes' theorem (see Section 4) to update their probabilities when new information arrives, see [52, Chapter 7] for a review of Savage's analysis in this context. In the next section we extend these ideas to a recent approach to decision theory and information theory.

3. Free energy

At this point we have yet to connect economic decisions with information theory. To move in that direction we introduce work that was developed using ideas from physics in such a way that statistical uncertainty in the choice an agent will make connects the *optimal* decisions described in the previous section with *sub-optimal* decisions and consequently information theory. The central idea from physics is that of the *free energy*.

3.1. Free energy in physics

We first introduce the Helmholtz free energy \mathcal{F}_p used in physics as the internal energy of the system $V(x)$ having discrete states x less the Shannon entropy of the system:

$$H(x) = - \sum_i P(x_i) \log(P(x_i)) \quad (3)$$

multiplied by the temperature of the system T :

$$\mathcal{F}_p = V(x) - TH(x). \quad (4)$$

This is the energy that is available to do work in a thermodynamic system at equilibrium. It removes the energy that is not available due to thermal fluctuations: the smaller the T value and consequently the smaller the thermal fluctuations are, the smaller the contribution from the entropy becomes, and so more of the total energy of the system V is available for doing thermodynamic work.

Equation 4 has an analogue in neuroscience that is the basis of what Friston calls the Free Energy Principle of the brain, but before describing that, we briefly introduce the same general principles in an economic context.

3.2. Free utility in economics

Equation 4 has an economic counterpart that appears, for example, in the earlier work of Wolpert [57,58] on predictive game theory and collective intelligence. In its simplest form an agent chooses a distribution p over the utilities $U(x_i)$ over a finite set of discrete choices $\{x_i\}$ such that they have an expected utility:

$$E_p[U(x)] = \sum_i p_i U(x_i) \quad (5)$$

The 'free utility' of this situation is given by:

$$\mathcal{F}_g = E_q[U(x)] - TH(x), \quad (6)$$

in which T is analogous to the thermodynamic temperature and represents a decision-maker's uncertainty in their choice: the lower their uncertainty (small T) the higher their free utility, and the entropy measures the utility lost due to the sub-optimality of their choice. In a game theory setting, Equation 6 can be derived from an exponential distribution of probabilities over choices that can be found by using Jaynes' method of maximising the entropy where the expected utility of the game is the constraint [13,14,59], this results in the well known quantal response equilibrium of McKelvey and Palfrey [60].

This is true in general where we might know, for example through experimental observation, that an economic problem has a constraint in the form of an expectation: $E_p[G(x)]$ and we would like to know the “least informative” [61] distribution that satisfies this constraint. Then Jaynes’ MaxEnt [62] provides a way of deriving the following distribution:

$$P(x_i) = \mathcal{Z}^{-1} e^{\beta E_p[G(x)]} \quad (7)$$

with \mathcal{Z} normalising the distribution and the free energy is:

$$\mathcal{F}(p) = \frac{\partial \mathcal{Z}}{\partial \beta} + \beta^{-1} \sum_x P(x) \log(P(x)), \quad (8)$$

$$\begin{aligned} &= E_p[G(x)] - \beta^{-1} H(x), \\ &= \text{Utility} - \beta^{-1} \text{Entropy}. \end{aligned} \quad (9)$$

If $\mathcal{F}(p)$ is then minimised over the distribution $p(x)$ we will arrive back at Equation 7.

A similar approach was taken in the work of Grünwald and Dawid on Bayes optimal decisions and game theory [63]. There they sought to generalise constrained maximum entropy methods by casting MaxEnt as minimising a log-loss function and then generalising the method to other loss functions. Amongst others, the following results emerged:

- Using MaxEnt to form constrained distributions over decision variables is compatible with Bayes optimal decisions,
- It brings maximising the entropy into correspondence with maximising the expected utility when the utility is a log-loss function,
- These methods can be generalised to finding other optimal, constrained strategies using a variety of loss functions.

These results illustrate that there is an underlying discipline-independent aspect to optimisation in decision theory that respects Savage’s Bayesian principles in economics and it can be framed in terms of entropic methods. In the next section we turn to the approach developed by Jaynes to model a robot decision-maker, a prototype of an idealised AI.

4. The logic of Jaynes’ learning robot

In *Probability Theory: The Logic of Science* [64] Jaynes introduces a robot (generalised AI) that acts as a rational decision-maker that follows certain desiderata that reflect rationality. For a history of the use of Bayes theory, going as far back as Herodotus and the necessity of reasoning, Jaynes has written an extensive review [65] that covers a broader context of learning, inference, and probability theory. For his own contributions in *Probability Theory*, he introduces his robot in the following way:

In order to direct attention to constructive things and away from controversial irrelevancies, we shall invent an imaginary being. Its brain is to be designed by us, so that it reasons according to certain definite rules. These rules will be deduced from simple desiderata which, it appears to us, would be desirable in human brains; i.e. we think that a rational person, should he discover that he was violating one of these desiderata, would wish to revise his thinking. [64, Chapter 1: The thinking computer]

The desiderata that are going to be built into the robot are in three parts [64, Chapter 1: Plausible reasoning]:

Desideratum (I):

Degrees of Plausibility are represented by real numbers. Desideratum (I) is practically forced on us by the requirement that the robot's brain must operate by the carrying out of some definite physical process. However, it will appear that it is also required theoretically; we do not see the possibility of any consistent theory without a property that is equivalent functionally to Desideratum (I).

Desideratum (II):

Qualitative Correspondence with common sense. We want to give our robot another desirable property for which honest people strive without always attaining; that it always reasons consistently.

Desideratum (IIIa):

If a conclusion can be reasoned out in more than one way, then every possible way must lead to the same result.

Desideratum (IIIb):

The robot always takes into account all of the evidence it has relevant to a question. It does not arbitrarily ignore some of the information, basing its conclusions only on what remains. In other words, the robot is completely non-ideological.

Desideratum (IIIc):

The robot always represents equivalent states of knowledge by equivalent plausibility assignments. That is, if in two problems the robot's state of knowledge is the same (except perhaps for the labeling of the propositions), then it must assign the same plausibilities in both.

Jaynes identifies these desiderata with how his rational robot should reason. Specifically, his robot uses the following fundamental principle to inform all probabilistic inferences that it makes:

To form a judgement about the likely truth or falsity of any proposition A , the correct procedure is to calculate the probability that A is true:

$$P(A|E_1E_2\ldots) \quad (10)$$

conditional on all the evidence at hand. [64, pg. 86]

Jaynes goes on to establish how the robot uses Bayes' theorem to learn and reason. He begins with a list of various propositions denoted A, B, C , etc., letting AB stand for the proposition "both A and B are true", \bar{A} for " A is false", and $P(A|B)$ represents the probability that A is true given B , Jaynes restates the product and sum rules of probability:

$$P(AB|C) = P(A|BC)P(B|C) \quad (11)$$

$$1 = P(A|B) + P(\bar{A}|B) \quad (12)$$

$$P(A|BC) = P(A|C) \frac{P(B|AC)}{P(B|C)} \quad (13)$$

where the last statement is Bayes' theorem. For Jaynes, the key point embedded in Bayes theorem is the functional ability for his robot to be able to learn from new information. To see this consider:

$$\begin{aligned}
 P(DH|X) &= P(H|DX)P(D|X) \\
 X &= \text{prior information} \\
 H &= \text{an hypothesis to test} \\
 D &= \text{the observed data} \\
 \Rightarrow P(H|DX) &= P(H|X) \frac{P(D|HX)}{P(D|X)}
 \end{aligned} \tag{14}$$

where the *prior probability* $P(H|X)$ is updated with new information D to give the *posterior probability* $P(H|DX)$, and Equation 13 provides Jaynes' reasoning robot with the computational steps to do so.

This is only the first part of the reasoning that the robot can do though. In [64, Chapter 11] Jaynes extends these foundations to how a robot might reason in a more versatile fashion, pointing out that the robot does not yet have a way of assigning probabilities to prior information at the outset of a problem. According to Jaynes initially all the robot knows how to do is to break the problem up into a set of exhaustive, mutually exclusive states and assign a uniform probability based on the principle of indifference. However, if there is some prior information that can be expressed as an expectation then there is a principled method by which the robot can use this information as *mean valued constraints* in order to establish a Bayesian prior from which the robot can begin to learn and reason. This is formally the same as the method Jaynes established in physics [62], i.e. it is precisely the maximum entropy method of statistical mechanics but expanded and applied to the robot's decision-making processes.

By now it has been established in the literature that Bayesian methods and maximum entropy methods are consistent with respect to each other and they both satisfy some criteria that various school's of thought agree a reasonable decision-maker should respect [63]. In particular, Grünwald and Dawid established the formal relationship between what had, prior to their work, been suspected but not entirely formalised: that there is "a close relationship between two problems that have customarily been regarded as distinct: that of maximizing entropy, and that of minimizing worst-case expected loss [...] these two problems are shown to be dual to each other, the solution to each providing that to the other."

With that in mind, there is a close and self-consistent relationship between maximising the free utility in game theory, maximising entropy subject to mean valued constraints, and Bayesian decision theory. In the following section some of these elements come together in the free energy principle that Friston has proposed for modelling human cognition in a naturalistic environment.

5. Friston's free energy principle

In order to introduce the free energy principle without the complexity of Friston's complete model, a simpler procedural framework is introduced first before introducing the complete distributions with more realistic inputs, this follows [66]. We start with an agent who knows that there is a joint distribution $P(s, o)$ that describes the relationship between sensory observations o and the real states of the world s (hidden to the agent) that generated them. This joint distribution can be decomposed into a conditional probability and a marginal probability: $P(o, s) = P(o|s)P(s)$, i.e. the Bayesian likelihood and prior, respectively. $P(o, s)$ is called the generative model and is described in more detail below. In Bayesian cognitive science, the generative model is said to be inverted to give

the posterior distribution $P(s|o) \propto P(o|s)P(s)$ also called the recognition model.

The difficulty with the direct application of this method is that the likelihood is often computationally expensive or even intractable for all but the simplest models. To address this we can estimate the posterior with an approximate posterior (described in detail below): $Q(s) \simeq P(s|o)$, i.e. the new recognition model $Q(s)$ is an approximation to the Bayesian posterior. In order to compute $Q(s)$ we minimise the free energy:

$$Q(s) = \arg \min_Q \mathcal{F}(Q) \Rightarrow Q(s) \simeq P(s|o) \quad (15)$$

$$\mathcal{F}(Q) = E_Q[\log(Q(s)) - \log(P(s|o))] \quad (16)$$

$$= E_Q[-\log(P(s|o))] - H(Q(s)) \quad (17)$$

which converts the problem of having to directly estimate the likelihood to a more tractable optimisation problem. This is functionally the same as the free utility of an expected log-loss function $-\log(P(s|o))$ where expectation is over $Q(s)$, minus the entropy of $Q(s)$.

Friston's wider framework expands on the structural form of the distributions just described. In order to do this we introduce the following discrete sets [67] where the *external world* is the environment outside of the agent's cognitive process, i.e. it is demarcated by the agent's sensory organs communicating states of the world to the agent's internal cognitive processes:

- A finite set of observations Ω about the external world.
- A finite set of actions A the agent can take in the external world.
- A finite set of hidden states S of the external world.
- A finite set of control states U that the agent can choose from in order to act in the external world.

With these sets we can describe the three probability distributions that are used in Friston's approach to describing an adaptive agent interacting with their environment.

Distribution 1: The first distribution is the *generative process* that produces an agent's observations $o_i \in \Omega$ over a discretely indexed period of time $t \in \{0, \dots, T\}$ such that $o(t) \in \Omega$ and $\tilde{o} = \{o(0), \dots, o(T)\}$ is a specific instance of a time series of $o(t)$ observations. This notation is used for the other state variables $s_i \in S$, $a_i \in A$, and $u_i \in U$ as well. With variables $s(t) \in S$ and $a(t) \in A$, the generative process is defined as the joint probability distribution:

$$R(\tilde{o}, \tilde{s}, \tilde{a}) = P(\{o(0), \dots, o(T)\} = \tilde{o}, \{s(0), \dots, s(T)\} = \tilde{s}, \{a(0), \dots, a(T)\} = \tilde{a}) \quad (18)$$

Distribution 2: The *generative model* over observations is a belief model that the agent has regarding how the external world generates the agent's observations. Given a time series of observations $o(t) \in \Omega$, of hidden states $s(t) \in S$, and of control states $u(t) \in U$ the generative model is:

$$\begin{aligned} P(\tilde{o}, \tilde{s}, \tilde{u} | m) &= P(\{o(0), \dots, o(T)\} = \tilde{o}, \{u(0), \dots, u(T)\} = \tilde{u}, \{s(0), \dots, s(T)\} = \tilde{s}) \\ &= P(\tilde{o} | \tilde{s}) P(\tilde{s}, \tilde{u} | \gamma, \tilde{a}) P(\gamma | m) \end{aligned} \quad (20)$$

in which $P(\tilde{o} | \tilde{s})$ is the likelihood of the sequence of observations \tilde{o} given the hidden sequence of states \tilde{s} and:

$$P(\tilde{s}, \tilde{u} | \gamma, \tilde{a}) = P(\tilde{u} | s(t), \gamma) P(s(t) | s(t-1), a(t-1)) \dots P(s(1) | s(0), a(0)) P(s(0) | a(0)) \quad (21)$$

is the agent's beliefs about the state transitions of the Markovian process that generates the next hidden state(s) of the world from the previous state and the agent's previous

action: $P(x(t)|x(t-1)s(t-1))$. $P(\gamma|m)$ is the likelihood of the parameters γ given the agent's model m of the external world [CHECK this interpretation]. The utility is derived from the first term in Equation 21:

$$\log(P(\tilde{u}|s(t), \gamma)) = \beta \mathbf{U}_{tot}(\tilde{u}, s(t)) \quad (22)$$

$$\mathbf{U}_{tot}(\tilde{u}, s(t)) = H(P(s(T)|s(t), \tilde{u})) + \sum -P(s(T)|s(t), \tilde{u}) \log(P(s(T)|m)) \quad (23)$$

In this expression the first term is the entropy and represents the utility associated with an agent's exploration of states and the second term is the expected utility of a log-loss function for the distribution $P(s(T)|m)$, which is just the expectation of the log-probability of the external world's final state $s(T)$ given the agent's priors encoding their goals m .

Distribution 3: The last distribution is the *approximate posterior* distribution over the hidden and control states:

$$Q(\tilde{s}, \tilde{u}) = P(\{s(0), \dots, s(T)\} = \tilde{s}, \{u(0), \dots, u(T)\} = \tilde{u}). \quad (24)$$

This is the approximate posterior distribution of Equation 15 that is needed to be found through the free energy optimisation process.

The free energy expression: Substituting the last two distributions into Equation 17 we have the following form:

$$\mathcal{F}(Q(\tilde{s}, \tilde{u})) = E_Q[-\log(P(\tilde{o}, \tilde{s}, \tilde{u} | m))] - H(Q(\tilde{s}, \tilde{u})), \quad (25)$$

$$\Rightarrow Q(\tilde{s}, \tilde{u}) = \arg \min_Q \mathcal{F}(Q(\tilde{s}, \tilde{u})). \quad (26)$$

This is the final form of Friston's free energy principle for how an agent estimates the posterior joint distribution of $Q(\tilde{s}, \tilde{u})$, i.e. by optimising in order to find the (approximate) Bayesian posterior distribution, the agent is implicitly maximising a utility function that is embedded in the second term of the generative model given by Equation 20. The utility is explicitly expressed in Equation 23, which in turn was derived from the first term of Equation 21. This is achieved by the agent by manipulating the control states \tilde{u} that maximise the expected utility of the final state of the external world $s(T)$.

6. Collecting the threads

What I have sought to illustrate with the discussion above are the three different perspectives each has adopted, and how they have been framed in different ways by different researchers, but the underlying formalisms are not very far apart. For example axiomatic decision theory has had a powerful influence on how we understand what we should want or expect from an ideal decision-maker, and once we think in terms of an 'ideal' these axioms are very appealing. But axiomatic methods do not necessarily reflect what happens under the cognitive hood of our psychology and so they can only be part of the solution to understanding how real agents make decisions, for example in new environmental contexts. Psychology and neuroscience take the opposite position to the axiomatic approach in saying that the process is the point at which investigation begins, and if a cohesive theory ultimately comes from that then all the better, but the empirical foundations are central. Finally, while AI looks under the hood to some extent, it also has a bi-modal approach: one mode emphasising the optimality of a data-intensive analysis, the other emphasising the implementation of some decision-making desiderata, not unlike economic decision-theory.

Each of these approaches can, in principle, be implemented in software or hardware, if only because they are describable mathematically, and in fact they were included here

for this very reason. This might not get us any closer to realising the true potential of these algorithms to model human-like decision-making for many decades due to a lack of computational power, but the potential may already be there. This would still leave us with the question of how much of our human abilities need to be retained in models of decision theory for a particular application, for example in a macro-economic model we might only need a very simple model of each agent but to realistically model two people playing a game like Go or chess we likely need to model most of a person's neuro-psychology such as perception, long term memory, and detailed strategising and mind-reading².

Before moving on to the main parts of this section it is important to note two particularly illustrative directions in artificial intelligence that have extended our understanding of what is theoretically and practically possible. Successfully playing Go against a professional player was considered at least a decade away when Lee Sedol was defeated in 2016, and that this was a grand challenge for AI, and since then the AlphaGo ecosystem of AIs has grown and been successful in other areas that were considered to be too challenging for current AI as well. However, getting to the moon cannot be achieved by technological breakthroughs in ladder design, and it is not clear whether our current AI technology is just a "ladder" or an under-powered "rocket" in terms of human intelligence and decision-making, and this debate has not yet played itself out by any means, see [49] for a wide-ranging view of the current limitations of AI. What is clear though is that there is likely to be a place for both rockets and ladders of many different designs in the future applications of AI. In the next subsections I discuss some potential directions for applying information theory to decision theory. In an alternative direction, Pearl has carried out extensive work on Bayesian networks and their applications to artificial intelligence and decision-theory [68,69]. The importance of this work lies in the potential for Bayesian belief networks to encode human-like expertise in a rigorous and implementable fashion. Of particular interest in the context of this review is that Pearl's work [70] has informed aspects of Da Costa and Friston's et al work on probabilistic graphical models for active inference [71]. Much like the results discussed above for artificial intelligence, this is very recent work and it looks to be a very promising direction for further work.

6.1. Don't look under the hood!

Expected utility theory has been the dominant paradigm in economic decision theory at least since the capitulation of Samuelson in 1950 [56], and despite recent challenges from alternatives such as behavioural economics [72], prospect theory [73], and maxmin expected utility theory [74] it still remains central to economic theory [56]. The theory is motivated by the principle that agents make choices *as if* they were maximising a utility function and that an economist can infer the existence of this function through the observation of an agent's preferences. In this frame the economist chooses not to look under the hood [75] of the agent making a decision, i.e. they deliberately choose not to understand psychological processes. This is a form of *substantive rationality*, *what* decisions are made is more important than *how* they are made [76]. In this setting an agent is rational if their *behaviour* is appropriate for the achievement of their goals (such as maximising an inferred utility function) within the conditions and constraints of their environment [77]. Simon summarises the point here:

Thus, the assumptions of utility or profit maximization, on the one hand, and the assumption of substantive rationality, on the other, freed economics from any dependence upon psychology. [77, pg. 66]

² In the psychological *theory of mind* sense.

Taking this view at face value, we ask: what can information theory contribute to economics in this mode? Recall from the Introduction that McGill proposed information theory as an instrument for investigating multivariate time series [20] where the patterns of observed behaviour of the agents are encoded in the data. These patterns, assuming that they are individually identified, can be treated as the *revealed preferences* of individual agents³. Now suppose we want to know how each agent's behaviour influences the others, a common question in game theory but also of interest more generally in studies of social and economic interactions. Expected utility theory provides no guidance on the nature of an agent's utility function and it might be linear or non-linear. Traditional methods such as the Pearson correlation can measure the linear correspondence between the patterns of behaviour of two agents, but if the patterns are nonlinear they can be missed entirely by conventional techniques while mutual information, for example, will measure both the linear and non-linear aspects of behavioural correspondences [78–80]. Furthermore, modern software can be extended to the analysis of discrete event data of a continuous time process with limited modifications [7], allowing researchers to infer temporally directed relationships between events, even in complex time series.

What we conclude is that, even if we choose not to look under the psychological hood, then information theory is a much more sensitive instrument for measuring non-linearities in behaviour, and therefore non-linear utilities, see for example [81].

6.2. OK, should we look under the hood anyway?

The alternative to *substantive rationality* is *procedural rationality* [77] in which we are concerned with *how* an agent makes a decision, rather than *what* decision they make. Hausman gives the following argument equivalent to Friedman's original *as if* logic in order to critique Friedman's *Positive Economics* [75, pg. 184]:

1. A good used car drives safely, economically, and comfortably. (over-simplified premise)
2. The only test of whether a used car is a good used car is whether it drives safely, economically, and comfortably. (invalidly from 1)
3. Anything one discovers by opening the hood and checking the separate components of a used car is irrelevant to its assessment. (trivially from 2)

So what is wrong with this argument? Hausman argues that fundamentally the future behaviour (performance) of the car cannot be reasonably, i.e. within cost and time bounds, assessed by simply taking it for a drive. Opening the hood and looking at the condition of the engine, looking at the chassis, and the current state of all the other components of the car will give a much clearer picture of what the car will cost for repair and maintenance into the future, whether or not it will breakdown in the next six months, and every used car buyer would be sensible to do such tests.

This analogy shares a great deal with another analogy, this one used by one of the founders of complexity economics, Brian Arthur [82], about the flow of traffic on a freeway⁴. In this analogy traffic can be modelled as a steady state equilibrium of traffic flows for much of the time, but traffic jams form and they do not stem from just one car or from every car, the jam is caused by a small number of cars and their interactions (for example the distance between them, breaking and accelerating times). So in order to understand the complex patterns that emerge in traffic flows we need to look beyond

³ There are some caveats that would need to be carefully considered, for example the complete set of choices available to each agent might not be known to the economist or the agent, and the environment and constraints might be non-stationary etc. but these are issues for any analysis of this type.

⁴ An interview: W. Brian Arthur on "Complexity Economics" (9th May 2015) <https://www.complexityexplorer.org/news/19-w-brian-arthur-on-complexity-economics>

equilibrium behaviour and understand the meso-scopic properties that cause these complexities to emerge.

Similarly for our decision-makers, if we choose to look under the psychological hood we need to find the right scale of analysis. Friston's free energy approach is interesting in this respect because it has both neuronal and psychological [83] explanatory power. In terms of what information theory may say about this approach, there is an important result that states, for a suitably defined utility function $U(\cdot)$ and any probability space (Ω, \mathcal{F}, P) we have that: $U(A|B) = f(P(A|B))$ and $f(\cdot) = \alpha \log(\cdot)$ ($\alpha > 0$) from which the conclusion follows that utility *is* information in the form of (scaled) surprisal, see [31] and references therein. While this is in agreement with the Friston model as well as the Grünwald and Dawid methods, it approaches the problem from a different direction and makes explicit that the log-loss function is the same as utility, and that expected utility is the expected log-loss.

So a key strength in looking under the hood is that we get a much richer view of the generative processes of behaviour than we would if all we knew was the current repertoire of behaviours in a static world of stationary preferences. This has an impact when for example new technologies become available or when an equilibrium state is disturbed, two facets of modern life that are becoming more and more prominent as our everyday experience is disrupted by events that are difficult to foresee based on forecasting from past events. In this way we can use simulations based on the underlying cognitive processes to test and make predictions about an agent's behaviour in circumstances that have not yet been observed.

6.3. How large is our world?

So far it might still seem plausible that a clever combination of 'big data', massive computing power, and some Bayesian brain modelling using neural networks might be enough to reasonably model human decision making, even if we need to wait a few decades for the necessary technology to make it practical. If this were true then, in terms of the theory (if not the hardware and software), we would essentially be done. This is the central concern of Binmore's book [52] as well as a central point of Savage's original thesis [53]: that Bayesianism is not always rational and when it is not rational is precisely where we fail to understand the difference between our theories of decision-making and how decision-making biological agents operate in the real world.

Savage [53, pg. 16] described a 'small world' as being the situation in which a decision-maker can always "look before you leap", i.e. all information about the consequences of each choice is available to the decision-maker. A 'large world' is the situation where "you can cross the bridge when you come to it", i.e. the consequences are not known before a decision is made, so the decision-maker will deal with them as they arise. For example game theory is a small world problem, but the financial markets are not. Binmore identifies the work of Gilboa and Schmeidler [84] as being a case in point where economics has attempted to tackle this issue by studying imitation as an aspect of rationality. This approach was the inspiration for an artificial neural network approach to expert-level understanding of the patterns in the game of Go [23,24], a situation that is very much in the 'large world' setting.

The issue that Savage and Binmore raise is both subtle and key to decision theory in each of the contexts covered earlier in this article. Binmore has recently written a summary of the problems he sees in the construction of Bayesian priors [85]:

There is no consensus on where such a logical prior comes from, but scholars who see Bayesianism as embodying the essence of rational learning

usually argue that we should place the choice of prior at a time in the past when Alice is completely ignorant. An appeal is then sometimes made to the “Harsanyi doctrine”, which says that different rational agents independently placed in a situation of complete ignorance will necessarily formulate the same common prior (Harsanyi 1977). But even if one accepts this doubtful principle, one is left wondering what this “rational prior” would look like.

In practice, some version of Laplace’s principle of insufficient reason is usually employed. Others take this position further by arguing that the complete ignorance assumption implies, for example, that the rational prior will maximize entropy (Jaynes and Bretthorst 2003). We are then as far from Savage’s view on constructing priors as it is possible to be. Instead of using all potentially available information in the small world to be studied in formulating a prior, one treats all such information as irrelevant.

The literature offers many examples that draw attention to the difficulties implicit in appeals to versions of Laplace’s principle. Similar objections can be made to any proposal for choosing a logical prior that depends only on the state space. Even in small worlds we seldom know the “correct” state space. In a large world, we have to update our state space as we go along—just as physicists are trying to do right now as they seek to create a theory of everything.

In brief, Bayesianism has no solid answer to the question of where logical priors come from. Attempts to apply versions of the principle of insufficient reason fail because they take for granted that the state space is somehow known a priori beyond any need for revision. The same goes for any proposal that makes the prior a function of the state space alone. But how can Alice be so certain about the state space when she is assumed to know nothing whatever about how likely different states may be?

As mentioned above, Jaynes answer to this is to say that prior information for his robot can be used by encoding it as an expectation and this, via MaxEnt, leads to a well defined statistical model that can then be updated through Bayesian learning. This still does not address the concerns Binmore raises, but it appears to take the debate one step closer to a practical resolution, if not a theoretical one that reflects an axiomatic approach to decision theory. Friston also addresses this issue for the free energy principle here:

Kersten et al. (2004) provide an excellent review of object perception as Bayesian inference and ask a fundamental question, ‘Where do the priors come from. Without direct input, how does image-independent knowledge of the world get put into the visual system?’ In §3(e), we answer this question and show how empirical Bayes allows priors to be learned and induced online during inference. [29]

So let us compare what Friston says in his §3(e) with what Binmore has to say about massaging information in order to get consistent priors, Friston’s central point is this:

Empirical Bayes allows these priors to be induced by sensory input, using hierarchies of backward and lateral projections that prevail in the real brain. In short, hierarchical models of representational learning are a natural choice for understanding real functional architectures and, critically, confer a necessary role on backward connections.

This pivots on the method of *empirical Bayes*, an approach to establishing priors directly from data in a hierarchical structure:

The problem with fully Bayesian inference is that the brain cannot construct the prior expectation and variability [...] *de novo*. They have to be learned and also adapted to the current experiential context. This is a solved problem in statistics and calls for empirical Bayes, in which priors are estimated from data. Empirical Bayes harnesses the hierarchical structure of a generative model, treating the estimates at one level as priors on the subordinate level (Efron & Morris 1973). This provides a natural framework within which to treat cortical hierarchies in the brain, each level providing constraints on the level below. This approach models the world as a hierarchy of systems where supraordinate causes induce and moderate changes in subordinate causes. These priors offer contextual guidance towards the most likely cause of the input.

In contrast, Binmore's appeal to 'massaging' the priors:

At the end of the story, the situation is as envisaged by Bayesianites: the final "massaged" posteriors can indeed be formally deduced from a final "massaged" prior using Bayes' rule. This conclusion is guaranteed by the use of a complex adjustment process that operates until consistency is achieved. [...]

Notice that it is not true in this story that [an agent] is learning when the massaged prior is updated to yield a massaged prior - that rational learning consists of no more than applying Bayes rule. On the contrary, Bayesian updating only takes place after all learning is over. [...] Bayesianites therefore have the cart before the horse. Insofar as learning consists of deducing one set of beliefs from another, it is the massaged prior that is deduced from the unmassaged posteriors. [86]

Binmore's first paragraph could be in a agreement with Friston's use of hierarchical methods from empirical Bayes, but the second paragraph makes it clear that Binmore does not see any room in 'Bayesianism' for empirical learning. What Binmore appears to be looking for in the second paragraph is the learning that occurs before any Bayesian updating or massaging occurs, but this is what Friston appears to be offering in his neural framework of a cortical hierarchy, although it is not clear from reading their respective views how either would respond without explicitly putting the case to each of them. Another intriguing point is that Friston's view appears to be compatible with a perceptual deep-learning AI with feed forward – feed back mechanisms, and in that sense the Large World of Savage and Binmore may well be (partially) addressed by these early sensory information processing mechanisms. This may well be an approach to understanding how the AlphaGo AI is able to perform so well in the Large(-ish) World of Go⁵, it is able to build priors dynamically and from scratch, further complicating the issue around the capabilities of AI's approach to theoretical limits in decision-making.

6.4. *Where to for the future?*

There are many reasons for exploring these different approaches and how they intersect with one another, two possibilities are mentioned here. The first is that AI could be used as a model of human reasoning to help us better understand and improve upon our socioeconomic systems [87]. We can see that some of our greatest risks are instabilities caused by interactions between opposing views in politics, economics, and other spheres of our lives, so can better models of how we make decisions individually

⁵ Savage and Binmore may consider Go to be a Large World, or at least larger than a Small World, but it is still very small relative to the real world and these AIs are still highly dependent on operating within the boundaries of their task environment

and collectively help fix these issues? This is a grand challenge but psychologically well-founded models of people's behaviour could help us understand how behaviour at the individual level causes patterns of instability at larger social scales. The second is that, as Brian Arthur has observed [82], the US (and the world) economy is becoming increasingly more autonomous, but can AI be used to supplement, rather than supplant, people in the workplace to make working conditions better? This needs highly sophisticated AIs that use accurate models of human psychology in order to interact with people effectively. However, the challenges are exceptional, AI has gone some small way to looking under the psychological hood and the small world of Savage has been made a little larger with recent advances on more complex tasks. But AIs are still prone to failure in only slightly modified environments that people would find trivial to navigate and they do not yet have the resilience and adaptability people have in making a decision. A particular hurdle is that, in each of the examples in this article, the agents do not appear to have a way to *understand* a decision-making task or to be able to work effectively when they cannot look before they leap. Melanie Mitchell, in *Artificial intelligence hits the barrier of meaning* [49] sums up the challenge for these areas in the future:

It may be that the pure information processing metaphor is not sufficient, and that other frameworks might give the insights that we need. Some proposals for such frameworks include free energy principles, control theory, dynamical systems theory, and ideas from biological development. It may be that an entirely new framework is needed for a full account of cognition.

Funding: This research received funding from the Australian Research Council, grant number DP170102927

Conflicts of Interest: The author declares no conflict of interest.

1. Georgescu-Roegen, N. The entropy law and the economic problem. *Valuing the Earth: Economics, ecology, ethics* **1993**, pp. 75–88.
2. Wilson, A. *Entropy in urban and regional modelling*; Vol. 1, Routledge, 2011.
3. Crosato, E.; Prokopenko, M.; Harré, M.S. The polycentric dynamics of Melbourne and Sydney: Suburb attractiveness divides a city at the home ownership level. *Proceedings of the Royal Society A* **2021**, *477*, 20200514.
4. Crosato, E.; Nigmatullin, R.; Prokopenko, M. On critical dynamics and thermodynamic efficiency of urban transformations. *Royal Society open science* **2018**, *5*, 180863.
5. Harré, M.; Bossomaier, T. Phase-transition-like behaviour of information measures in financial markets. *EPL (Europhysics Letters)* **2009**, *87*, 18009.
6. Bossomaier, T.; Barnett, L.; Steen, A.; Harré, M.; d'Alessandro, S.; Duncan, R. Information flow around stock market collapse. *Accounting & Finance* **2018**, *58*, 45–58.
7. Harré, M. Entropy and Transfer Entropy: The Dow Jones and the Build Up to the 1997 Asian Crisis. *Proceedings of the International Conference on Social Modeling and Simulation, plus Econophysics Colloquium 2014*. Springer, Cham, 2015, pp. 15–25.
8. Bossomaier, T.; Barnett, L.; Harré, M. Information and phase transitions in socio-economic systems. *Complex Adaptive Systems Modeling* **2013**, *1*, 1–25.
9. Matsuda, H.; Kudo, K.; Nakamura, R.; Yamakawa, O.; Murata, T. Mutual information of Ising systems. *International Journal of Theoretical Physics* **1996**, *35*, 839–845.
10. Barnett, L.; Lizier, J.T.; Harré, M.; Seth, A.K.; Bossomaier, T. Information flow in a kinetic Ising model peaks in the disordered phase. *Physical review letters* **2013**, *111*, 177203.
11. Prokopenko, M.; Barnett, L.; Harré, M.; Lizier, J.T.; Obst, O.; Wang, X.R. Fisher transfer entropy: quantifying the gain in transient sensitivity. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **2015**, *471*, 20150610.
12. Aoki, M.; others. *New Approaches to Macroeconomic Modeling*. Cambridge Books **1998**.
13. Wolpert, D.H.; Harré, M.; Olbrich, E.; Bertschinger, N.; Jost, J. Hysteresis effects of changing the parameters of noncooperative games. *Physical Review E* **2012**, *85*, 036102.
14. Harré, M.S.; Bossomaier, T. Strategic islands in economic games: Isolating economies from better outcomes. *Entropy* **2014**, *16*, 5102–5121.
15. Lizier, J.T. JIDT: An information-theoretic toolkit for studying the dynamics of complex systems. *Frontiers in Robotics and AI* **2014**, *1*, 11.
16. Bossomaier, T.; Barnett, L.; Harré, M.; Lizier, J.T. An introduction to transfer entropy. *Cham: Springer International Publishing* **2016**, 65.

17. Laming, D. Statistical information and uncertainty: A critique of applications in experimental psychology. *Entropy* **2010**, *12*, 720–771.
18. Shannon, C.E. A mathematical theory of communication. *The Bell system technical journal* **1948**, *27*, 379–423.
19. Miller, G.A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* **1956**, *63*, 81–97.
20. McGill, W. Multivariate information transmission. *Transactions of the IRE Professional Group on Information Theory* **1954**, *4*, 93–111.
21. Luce, R.D. Whatever happened to information theory in psychology? *Review of general psychology* **2003**, *7*, 183–188.
22. Harré, M.S.; Bossomaier, T.; Gillett, A.; Snyder, A. The aggregate complexity of decisions in the game of Go. *The European Physical Journal B* **2011**, *80*, 555–563.
23. Harré, M.; Bossomaier, T.; Snyder, A. The development of human expertise in a complex environment. *Minds and Machines* **2011**, *21*, 449–464.
24. Harré, M.; Snyder, A. Intuitive expertise and perceptual templates. *Minds and Machines* **2012**, *22*, 167–182.
25. Wollstadt, P.; Lizier, J.T.; Vicente, R.; Finn, C.; Martinez-Zarzuela, M.; Mediano, P.; Novelli, L.; Wibral, M. IDTxL: The Information Dynamics Toolkit xl: a Python package for the efficient analysis of multivariate information dynamics in networks. *The Journal of Open Source Software* **2019**. doi:10.21105/joss.01081.
26. Wibral, M.; Vicente, R.; Lindner, M. Transfer entropy in neuroscience. In *Directed information measures in neuroscience*; Springer, 2014; pp. 3–36.
27. Wibral, M.; Lizier, J.; Vögler, S.; Priesemann, V.; Galuske, R. Local active information storage as a tool to understand distributed neural information processing. *Frontiers in neuroinformatics* **2014**, *8*, 1.
28. Cramer, B.; Stöckel, D.; Kreft, M.; Wibral, M.; Schemmel, J.; Meier, K.; Priesemann, V. Control of criticality and computation in spiking neuromorphic networks with plasticity. *Nature communications* **2020**, *11*, 1–11.
29. Friston, K. A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences* **2005**, *360*, 815–836.
30. Genewein, T.; Leibfried, F.; Grau-Moya, J.; Braun, D.A. Bounded rationality, abstraction, and hierarchical decision-making: An information-theoretic optimality principle. *Frontiers in Robotics and AI* **2015**, *2*, 27.
31. Braun, D.A.; Ortega, P.A.; Theodorou, E.; Schaal, S. Path integral control and bounded rationality. 2011 IEEE symposium on adaptive dynamic programming and reinforcement learning (ADPRL). IEEE, 2011, pp. 202–209.
32. Ortega, P.A.; Braun, D.A. Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **2013**, *469*, 20120683.
33. Friston, K.J.; Daunizeau, J.; Kiebel, S.J. Reinforcement learning or active inference? *PloS one* **2009**, *4*, e6421.
34. Friston, K.; Schwartenbeck, P.; FitzGerald, T.; Moutoussis, M.; Behrens, T.; Dolan, R.J. The anatomy of choice: active inference and agency. *Frontiers in human neuroscience* **2013**, *7*, 598.
35. Bishop, C.M. *Pattern recognition and machine learning*; springer, 2006.
36. Friston, K.; Kilner, J.; Harrison, L. A free energy principle for the brain. *Journal of Physiology-Paris* **2006**, *100*, 70–87.
37. Oizumi, M.; Albantakis, L.; Tononi, G. From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. *PLoS Comput Biol* **2014**, *10*, e1003588.
38. Veness, J.; Ng, K.S.; Hutter, M.; Uther, W.; Silver, D. A monte-carlo ai approximation. *Journal of Artificial Intelligence Research* **2011**, *40*, 95–142.
39. Bolon-Canedo, V.; Remeseiro, B. Feature selection in image analysis: a survey. *Artificial Intelligence Review* **2019**, pp. 1–27.
40. Solorio-Fernández, S.; Carrasco-Ochoa, J.A.; Martínez-Trinidad, J.F. A review of unsupervised feature selection methods. *Artificial Intelligence Review* **2020**, *53*, 907–948.
41. Hall, M.A.; Smith, L.A. Practical feature subset selection for machine learning **1998**.
42. Huang, H.; Huang, J.; Feng, Y.; Zhang, J.; Liu, Z.; Wang, Q.; Chen, L. On the improvement of reinforcement active learning with the involvement of cross entropy to address one-shot learning problem. *PloS one* **2019**, *14*, e0217408.
43. Kamatani, N. Genes, the brain, and artificial intelligence in evolution. *Journal of Human Genetics* **2020**, pp. 1–7.
44. Ferrucci, D.A. Introduction to “this is watson”. *IBM Journal of Research and Development* **2012**, *56*, 1–1.
45. Silver, D.; Huang, A.; Maddison, C.J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; others. Mastering the game of Go with deep neural networks and tree search. *nature* **2016**, *529*, 484–489.
46. Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; others. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **2018**, *362*, 1140–1144.
47. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; others. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* **2020**.
48. Mitchell, M. Can GPT-3 Make Analogies? *Medium* **2020**.
49. Mitchell, M. Artificial intelligence hits the barrier of meaning. *Information* **2019**, *10*, 51.
50. Damielson, P.; Audi, R.; Bicchieri, C.; others. *The Oxford handbook of rationality*; Oxford University Press, USA, 2004.
51. Steele, K.; Stefánsson, H.O. Decision Theory. In *The Stanford Encyclopedia of Philosophy*, Winter 2020 ed.; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University, 2020.
52. Binmore, K. *Rational decisions*; Princeton University Press, 2008.
53. Savage, L.J. *The foundations of statistics*; Courier Corporation, 1972.
54. Friedman, M. The methodology of positive economics **1953**.

55. Maki, U.; Mäki, U. *The methodology of positive economics: Reflections on the Milton Friedman legacy*; Cambridge University Press, 2009.
56. Moscati, I. Retrospectives: how economists came to accept expected utility theory: the case of samuelson and savage. *Journal of economic perspectives* **2016**, *30*, 219–36.
57. Wolpert, D.H. Predictive game theory **2005**.
58. Wolpert, D.H.; Wheeler, K.R.; Tumer, K. Collective intelligence for control of distributed dynamical systems. *EPL (Europhysics Letters)* **2000**, *49*, 708.
59. Harré, M.S.; Atkinson, S.R.; Hossain, L. Simple nonlinear systems and navigating catastrophes. *The European Physical Journal B* **2013**, *86*, 289.
60. McKelvey, R.D.; Palfrey, T.R. Quantal response equilibria for normal form games. *Games and economic behavior* **1995**, *10*, 6–38.
61. Niven, R.K. Jaynes' MaxEnt, steady state flow systems and the maximum entropy production principle. AIP Conference Proceedings. American Institute of Physics, 2009, Vol. 1193, pp. 397–404.
62. Jaynes, E.T. Information theory and statistical mechanics. *Physical review* **1957**, *106*, 620.
63. Grünwald, P.D.; Dawid, A.P.; others. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *the Annals of Statistics* **2004**, *32*, 1367–1433.
64. Jaynes, E.T. *Probability theory: The logic of science*; Cambridge university press, 2003.
65. Jaynes, E.T. Bayesian methods: General background **1986**.
66. Ramstead, M.J.; Kirchhoff, M.D.; Friston, K.J. A tale of two densities: Active inference is enactive inference. *Adaptive Behavior* **2020**, *28*, 225–239.
67. Friston, K.; Rigoli, F.; Ognibene, D.; Mathys, C.; Fitzgerald, T.; Pezzulo, G. Active inference and epistemic value. *Cognitive neuroscience* **2015**, *6*, 187–214.
68. Pearl, J. *Probabilistic reasoning in intelligent systems: networks of plausible inference*; Elsevier, 2014.
69. Pearl, J. Causes of effects and effects of causes. *Sociological Methods & Research* **2015**, *44*, 149–164.
70. Pearl, J. Graphical models for probabilistic and causal reasoning. *Quantified representation of uncertainty and imprecision* **1998**, pp. 367–389.
71. Da Costa, L.; Parr, T.; Sajid, N.; Veselic, S.; Neacsu, V.; Friston, K. Active inference on discrete state-spaces: a synthesis. *Journal of Mathematical Psychology* **2020**, *99*, 102447.
72. Mullainathan, S.; Thaler, R.H. Behavioral economics. Technical report, National Bureau of Economic Research, 2000.
73. Tversky, A.; Kahneman, D. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* **1992**, *5*, 297–323.
74. Gilboa, I.; Schmeidler, D. Maxmin expected utility with non-unique prior. *Journal of mathematical economics* **1989**, *18*, 141–153.
75. Hausman, D.M. Why look under the hood. *The philosophy of economics: An anthology* **1994**, pp. 217–221.
76. Binmore, K. Rationality. In *Handbook of game theory with economic applications*; Elsevier, 2015; Vol. 4, pp. 1–26.
77. Simon, H.A. From substantive to procedural rationality. In *25 years of economic theory*; Springer, 1976; pp. 65–86.
78. Li, W. Mutual information functions versus correlation functions. *Journal of statistical physics* **1990**, *60*, 823–837.
79. Cellucci, C.J.; Albano, A.M.; Rapp, P.E. Statistical validation of mutual information calculations: Comparison of alternative numerical algorithms. *Physical Review E* **2005**, *71*, 066208.
80. Wang, Q.; Shen, Y.; Zhang, J.Q. A nonlinear correlation measure for multivariable data set. *Physica D: Nonlinear Phenomena* **2005**, *200*, 287–295.
81. Harré, M.S. Strategic information processing from behavioural data in iterated games. *Entropy* **2018**, *20*, 27.
82. Arthur, W.B. Foundations of complexity economics. *Nature Reviews Physics* **2021**, pp. 1–10.
83. Moran, R.J.; Campo, P.; Symmonds, M.; Stephan, K.E.; Dolan, R.J.; Friston, K.J. Free energy, precision and learning: the role of cholinergic neuromodulation. *Journal of Neuroscience* **2013**, *33*, 8227–8236.
84. Gilboa, I.; Schmeidler, D. *A theory of case-based decisions*; Cambridge University Press, 2001.
85. Binmore, K. On the foundations of decision theory. *Homo Oeconomicus* **2017**, *34*, 259–273.
86. Binmore, K. Rational decisions in large worlds. *Annales d'Economie et de Statistique* **2007**, pp. 25–41.
87. Parkes, D.C.; Wellman, M.P. Economic reasoning and artificial intelligence. *Science* **2015**, *349*, 267–272.