DEPARTMENT OF INTERNATIONAL AND EUROPEAN ECONOMIC STUDIES

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Rationality in Economics: Epistemic Assumptions and Pragmatic Justifications

Phoebe Koundouri

Nikitas Pittis

Panagiotis Samartzis

Working Paper Series

25-01

January 2025

Rationality in Economics: Epistemic Assumptions and Pragmatic Justifications"

Phoebe Koundouri*† Nikitas Pittis‡ Panagiotis Samartzis*§

Abstract

This paper has three main objectives. First, it aims to clarify the alternative concepts of rationality used in economics, specifically Rationality-A (coherence of subjective probabilities for every point in time), Rationality-B (perfect alignment with objective probabilities for every point in time), and Rationality-BB (asymptotic convergence to objective probabilities over time). Second, it seeks to identify distinct sets of epistemic assumptions (S1–S5) that logically entail each of these three definitions of rationality and to explain their respective roles. Third, it evaluates the pragmatic justifications for these assumptions, focusing on Dutch Book arguments for coherence and arbitrage arguments for asymptotic accuracy.

Keywords: Rationality, Coherence, Accuracy, Epistemic Assumptions, Pragmatic Justifications

JEL Classification: C44, D81, D83, D89

*These authors acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 951424; ERC-SyG Water-Futures).

†School of Economics and ReSEES Laboratory, Athens University of Economics and Business; Sustainable Development Unit, Athena Research Center; World Academy of Arts and Science; UN SDSN-Europe; e-mail: pkoundouri@aueb.gr, phoebe.koundouri@icre8.eu (corresponding author)

‡Department of Banking and Financial Management, University of Piraeus, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: npittis@unipi.gr

§Department of Banking and Financial Management, University of Piraeus, Greece; School of Economics, Athens University of Economics and Business, Greece; International Center for Research on the Environment and the Economy, Greece; e-mail: psamartzis@unipi.gr

1 Introduction

The concept of rationality plays a foundational role in economics and other disciplines such as philosophy, psychology, and decision theory. Rationality can be understood in terms of its object: preferences or beliefs. Preferences pertain to how an agent ranks alternatives, while beliefs pertain to how an agent forms subjective probabilities about uncertain events. Rationality seeks to impose structure and consistency on these objects, leading to various definitions and frameworks.

The most basic notion of rationality, referred to as Rationality-0, is grounded in preferences: An agent is said to be rational if her system of preferences satisfies the axioms of a formal preference calculus, such as Savage's (1974). These axioms, including completeness, transitivity, and independence, ensure that the agent's preferences can be represented as if she maximizes a utility function.

Alternatively, rationality refers to beliefs, with two primary definitions distinguished by the property of beliefs they emphasize. Rationality-A focuses on coherence, requiring that an agent's subjective probabilities obey the rules of probability calculus. Rationality-B, in contrast, emphasizes accuracy, mandating that an agent's subjective probabilities align with the objective probabilities derived from the true data-generating process.

The primary objectives of this paper are threefold. First, it aims to provide precise definitions of Rationality-A and Rationality-B, establishing clear standards for coherence and accuracy in belief formation. Second, it investigates the sets of assumptions under which these definitions hold, identifying the conditions - such as prior coherence, Bayesian updating, and principles of inference - that logically entail Rationality-A and Rationality-B. Third, it examines the justifications for these assumptions, focusing on the cognitive and epistemic properties that an agent must possess to satisfy them and assessing their feasibility and relevance to real-world decision-making.

2 Rationality-A (Coherence) and Rationality-B (Accuracy)

Rationality seeks to impose structure on how agents form and update beliefs. As already mentioned, the two primary definitions of rationality, Rationality-A and Rationality-B, focus on different properties of beliefs: coherence and accuracy, respectively. In what follows, Ω represents the sample space of all possible outcomes, and $\mathcal{F}$ is a $\sigma-$algebra of events (or propositions), which are subsets of Ω. $\mathcal{F}$ serves as the domain of the agent's subjective probability function(s).

Definition 1 *An agent is said to be Rational-A if her subjective probabilities, $P_t(A)$, $A \in \mathcal{F}$ at each point in time t obey the rules of probability calculus.*

This definition ensures that the agent's beliefs are internally consistent and coherent, adhering to principles such as non-negativity, normalization, and additivity.

Definition 2 *An agent is said to be Rational-B if her subjective probabilities, $P_t(A)$, $A \in \mathcal{F}$ at each point in time t, $t > 0$, are equal to the corresponding objective probabilities, $Ch(A)$.*

A related definition, referred to as Rationality-BB, introduces the asymptotic version of Rationality-B:

Definition 3 *An agent is Rational-BB if her subjective probabilities $P_t(A)$, $A \in \mathcal{F}$ converge to the corresponding objective probabilities $Ch(A)$ asymptotically.*

Rationality-BB relaxes the strong rationality constraint of Rationality-B allowing for a gradual alignment between subjective and objective probabilities.

Next, consider the following five key assumptions, which will be used to justify the aforementioned definitions of rationality. Each definition will be justified based on a distinct subset of these assumptions, as will be specified in the sequel.

Assumption S1: Coherence of the Prior

The agent's initial probability function $P_0 : \mathcal{F} \to [0.1]$ satisfies the axioms of probability calculus:

$$P_0(A) \geq 0,\ P_0(\Omega) = 1,\ P_0(A \cup B) = P_0(A) + P_0(B) \text{ for disjoint } A, B \in \mathcal{F} \quad (1)$$

Assumtion S2: Bayesian Conditionalization

The agent updates her subjective probabilities at time t, $t > 0$ via Bayesian Conditionalization (BC). If E_t represents the evidence observed up to and including time t, then the agent's posterior probability at t for any $A \in \mathcal{F}$ is given by

$$P_t(A) = P_0(A \mid E_t) \quad (2)$$

where

$$P_0(A \mid E_t) = \frac{P_0(A \cap E_t)}{P_0(E_t)},\ P_0(E_t) > 0. \quad (3)$$

Bayesian conditionalization ensures the dynamic consistency of an agent's beliefs over time. At $t = 0$, the agent effectively makes a self-commitment: "If, at a future time t I observe evidence E_t I shall set my subjective probability $P_t(A)$ of A equal to my prior conditional probability $P_0(A \mid E_t)$. Similarly, if I observe E'_t instead, I shall set $P_t(A) = P_0(A \mid E'_t)$. This commitment guarantees that the agent's posterior beliefs remain aligned with her prior conditional beliefs, avoiding dynamic incoherence. It is worth noticing that the agent determines these prior conditional probabilities $P_0(A \mid E_t)$ for every possible piece of evidence E_t only once in her lifetime, at $t = 0$. For the remainder of her epistemic life, her belief-updating process involves only two steps: (1) observing which piece of evidence E_t has occurred and (2) recollecting her initial self-commitment of how to update her beliefs if that evidence were ever the case.

Assumption S3: Principle of Direct Inference (Principal Principle)

The Principle of Direct Inference (PDI) refers to a class of rules connecting objective probabilities (or chances) to an agent's subjective probabilities. These rules guide how an agent should align her beliefs with objective probabilities when such probabilities are available. Among these principles, the Principal Principle (PP), proposed by David Lewis (1980), is perhaps the most prominent. PP states that if the agent knows that the objective probability $Ch(A)$ of an event A is equal to x then her prior probability P_0 should satisfy the following "coordination" condition:

$$P_0(A \mid Ch(A) = x,\ I) = x, \tag{4}$$

where I represents admissible information for A. Specifically, I is admissible for A if it contains no information for A beyond the information contained in the objective probability of A (see Lewis 1980). The following remarks are in order:

(i) The Principal Principle (PP) sits at the intersection of two competing interpretations of probability: subjective probability, understood as an agent's rational degrees of belief, and objective probability, understood as an external feature of the world. At first glance, PP seems to prioritize objective probability, as it explicitly requires an agent to defer to objective chances when they are known. This deference suggests a view where objective probabilities sit in the driver's seat, guiding subjective probabilities. On this view an agent who adheres to PP is called Objective Bayesian.

(ii) However, Lewis, who introduced PP, seems to disagree (see Lewis, 1980). For him, subjective probability is the more fundamental notion, with objective probabilities emerging as secondary constructs. Without subjective probabilities, objective chances would have no clear role - they would just exist as abstract statistical facts about the world, disconnected from belief or action. For example, knowing that the objective chance of rain tomorrow is 0,7 only matters because it influences how a rational agent should set her subjective probability of rain. Hence, on the Lewisian view, the primary role of objective probabilities is to serve as constraints on an agent's subjective degrees of belief through principles like PP. To borrow an analogy, subjective probabilities are the engine of rational reasoning, while objective probabilities are the fuel that powers this engine.

(iii) Technically, PP represents a constraint on agent's prior probability function P_0. In this sense, it is atemporal, since it does not prescribe how probabilities evolve over time - unlike BC, which explicitly addresses dynamic belief updating. In other words, while PP constrains the initial system of subjective probabilities, BC governs their temporal evolution.

Assumption S4: Agent Knows the True Hypothesis H

The agent knows the true hypothesis H which specifies the probabilistic properties of the data-generating process and entails the objective chance function $Ch(\cdot)$. For example, the agent knows that the coin tosses are represented by a sequence of independent and identically distributed (i.i.d) random variables, each following a Bernoulli distribution with parameter $\theta = 0.5$. This hypothesis H implies that the chance of any single toss landing Heads is $Ch(A) = 0.5$.

Assumption S5: Hypothesis Space Includes the True Hypothesis

The agent conceives a hypothesis space $\mathcal{H} = \{H_1, H_2, ..., H_n\}$ where each hypothesis H_i represents a possible data-generating process and entails a distinct chance function $Ch_i(\cdot)$. These hypotheses are assumed to be mutually exclusive (only one can be true) and collectively exhaustive (one of them must be true). In this sense, the hypothesis space $\mathcal{H}$ can be viewed as a partition of Ω. More specifically, it is assumed that the true hypothesis H - which specifies the probabilistic properties of the data-generating process and entails the true chance function $Ch(\cdot)$ - is included in this partition.

Now we are ready to state the followin three propositions:

Proposition 1: *Assumptions S1 and S2 entail Rationality-A.*

Proof: See Appendix.

Next, we consider the second proposition:

Proposition 2: *Assumptions S1, S3 and S4 entail Rationality-B.*

Assumption S3 establishes PP, which ensures that the agent is willing to defer to objective chances, provided they are known. Since PP pre-supposes the existence of a prior subjective probability function, S1 is also required. However, PP alone does not guarantee that the agent actually knows the objective chances it merely sets up a willingness to align with them when they are available. This is where S4 becomes essential. Assumption S4 states that the agent knows the true hypothesis H, which specifies the probabilistic properties of the data-generating process and entails the true chance function $Ch(\cdot)$. As a result, the agent is in possession of the objective chances for every event $A \in \mathcal{F}$. Together, S3 and S4 ensure that the agent's subjective probabilities align with the objective chances.

However, a subtle issue arises: When does the agent come to know H**?** At $t = 0$, the agent is assumed to be epistemically a *tabula rasa*. Therefore, the knowledge of H must occur at some later point in time, say $t = 1$. This explains why Rationality-B is defined with the restriction $t > 0$.

FInally we present our third proposition:

Proposition 3: *Assumptions S1, S2, S3 and S5 entail Rationality-BB.*

Proof: See Appendix.

3 Justifications and Realism of Assumptions S1, S2, S3, S4 and S5

We first examine arguments for the plausibility and realism of S1.

3.1 Coherence of the Prior (S1): Dutch Book Arguments

There are two main arguments for S1, hereafter referred to as AR1 and AR2.[1] AR1 utilizes the so-called "representation theorems", according to which S1 follows logically from a set of plausible rationality constraints (e.g. transitivity) on the agent's preference relation $\succeq$: P_0 is coherent iff $\succeq$ satisfies (for example) Savage's (1974) axioms. In this type of argument, "the burden of proof" has moved from beliefs to preferences. But then the question that was originally raised for beliefs now re-emerges for preferences: What are the arguments (if any) supporting the view that rationality of preferences amounts to $\succeq$ satisfying the axioms of Savage? The main such argument is the so-called "Money Pump Argument" (MPA) which is structurally similar to AR2, to which we now turn.

AR2 is the (in)famous Dutch Book Argument (DBA). DBA is based on the so-called Dutch Book theorem, which shows that an agent, X, who does not have coherent beliefs is susceptible to a Dutch Book. The latter is a set of bets, each of which appears to be fair to X (by her own standards) but all together assure that X will incur a net financial loss, whatever the outcome turns out to be. Since a rational X does not exhibit this type of susceptibility (so the argument goes) it follows logically (by a modus tollens argument) that she entertains coherent beliefs[2].

A typical DBA has the following form (adapted from Hajek 2008) Note that in the following discussion we will refer to $A \in \mathcal{F}$ as a proposition rather than an event::

P1: X's subjective probability of A, $A \in \mathcal{F}$, is given by X's betting price on the truth of the proposition A.

P2: (Dutch Book Theorem - DBT): If X's system of betting prices on the propositions of $\mathcal{F}$ violate the rules of probability calculus, then there is a Dutch Book consisting of bets at those prices.

P3: If there is a Dutch Book consisting of bets at X's betting prices, then X is susceptible to financial losses.

P4: If X is susceptible to financial losses, then she is irrational.

C(onclusion): If X's subjective probabilities violate the rules of probability calculus, then X is irrational.

This argument is valid (by "transitivity of implication and modus ponens). Whether it is sound is another matter. Let us for the moment assume that the premises **P1-P4** are true (the argument is sound) and clarify what DBA has established and what has not. The basic result of DBA is proposition **C**. From this it immediately follows that the proposition **Ca**: "If X is rational then her subjective probabilities do not violate the rules of probability calculus" is true

[1] Joyce (1998, 2009) offers a third argument for Probabilism, namely the "Accuracy Domination Argument". This argument (being "epistemic rather than pragmatic) is based on the idea that a rational DM prefers to form beliefs that are as close to the true ones as possible. Joyce's s main result is that under a set of axioms, if DM's probability function P does not obey the rules of probability calculus, then there exists another probability function P' which obeys these rules and is closer to the truth in every possible world.

[2] DBAs were introduced by Ramsey (1926) and de Finetti (1937). For a recent survey of these arguments see Hajek (2008).

(contrapositive of C). Similarly, the propositions **Cb**: "If X is not susceptible to the financial losses (as described above), then her subjective probabilities do not violate the rules of probability calculus" (contrapositive of the conclusion from **P2** and **P3**) and **Cc**: "If there is no Dutch Book consisting of bets at X's betting prices, then X's system of betting prices (subjective probabilities) do not violate the rules of probability calculus" are also true (contrapositive of **P2**).

However, DBA is silent on the truth of the proposition **P2a**: "If X's system of betting prices (subjective probabilities) do not violate the rules of probability calculus, then there is no Dutch Book consisting of bets at those prices". The truth of this proposition, usually referred to as the Converse Dutch Book Theorem (CDBT) was proved by Lehman (1955) and Kemeny (1955). DBT and CDBT may be joined together as follows: "X's system of betting prices on the propositions of $\mathcal{F}$ does not violate the rules of probability calculus, *if and only if* there is no Dutch Book consisting of bets at those prices.

It is worth emphasizing that in the context of both DBA and MPA, X's aversion to certain financial losses forces her to obey a calculus (namely, Kolmogorov's or Savage's calculus) that she may not know or might have never heard of. On this view, what DBA and MPA have (allegedly) managed to accomplish is to translate a "practical" or "pragmatic" notion of rationality into a "formal" one. In other words, X's attempts to avoid being money-pumped or Dutch-booked make her develop a system of "thought police" that "clubs her into line when she violates certain principles of right reasoning" (Garber 1983, p. 101). This means that both DBA and MPA are based on the fundamental assumption that DM is capable to implement a heuristic error-correction process, ECC, that yields rational beliefs and rational preferences. For the case of beliefs, ECC may be outlined as follows: X begins with a probability function, P_0, and examines whether, under P_0, she is susceptible to a Dutch Book. If she is, then she denies that Book and adjusts her initial P_0 (corrects her initial error) until her susceptibility disappears. This in turn implies that ECC may be thought of as a rational adaptation processes for achieving some specific goals which, by its very nature, applies to any decision maker, regardless of X's level of expertise in the relevant subject matter. In other words, is ECC is a purely a-priori process, and as such it does not depend on the presence of any empirical information and X's ability to process this information. Besides, the very conception of DBA was motivated by the ambition to make formal rationality accessible to the ordinary decision maker. To sum up, the only properties that an ordinary X is required to have in order to reach rationality are, first, an aversion to suffering a sure monetary loss and, second, the analytical skills to implement ECC.

Realism of S1

The realism of S1 hinges critically on whether the typical agent X is capable of implementing the aforementioned ECC - the heuristic procedure through which an agent detects and corrects incoherences in her probabilistic beliefs. This ability is central to the DBA, as the argument implicitly assumes that a rational agent not only *wishes to avoid* guaranteed financial loss but also *possesses the cognitive capacity to identify* potential incoherences in her belief system. More specifically, X's ability to perform ECC is hidden in Premise **P4**, which asserts that susceptibility to a Dutch Book implies irrationality. This premise relies on an implicit assumption, **P4a**, which states: "X is capable of implementing ECC across each and every Dutch Book that can be made against her." In other words, X must be capable of executing ECC across all conceivable Dutch Books, regardless of their complexity.

Is it realistic to assume that X has the skills to carry out the necessary calculations that are involved in ECC? The answer to this question depends on the degree of complexity of the Dutch Book that X is faced with. For simple Dutch Books, X is likely to have the computational capacity to detect the aforementioned incoherence. However, in order for X to be deemed as rational, she must be able to repel *any conceivable* Dutch Book made against her, simple or complex. A complex Dutch book may be thought of as one that requires very complex calculations before its financial consequences are deduced. The degree of complexity of a Dutch Book depends on the number of rows in the relevant pay-off matrix. This number increases with the number of propositions (bets), on which the Dutch Book is based, at an exponential rate. As a result, relatively simple Dutch Books (that is, those based on a small number of propositions) exhibit a high degree of complexity, which in turn renders ECC infeasible. Put differently, Dutch Books that appear to be "simple" in terms of the number of propositions involved may be "complex" in terms of computational tractability, "given the severe time and memory limitations of a 'fast and frugal' cognitive system." (Oaksford and Chater, 2007 p. 16). It seems that as we move up in the scale of complexity, DBAs run out of steam. In order to execute ECC for every Dutch book, DM must possess "computational omnipotence".

To sum up, if the agent possesses the ability to detect incoherence across all Dutch Books but fails to exercise it, this failure can reasonably be seen as a sign of irrationality. On the other hand, if the agent genuinely lacks this computational ability - not out of negligence or unwillingness, but because the problem exceeds human cognitive capacity - then susceptibility to incoherence reflects a computational limitation rather than irrationality.

The implicit requirement of agent's "computational omnipotence" discussed above can be seen as a practical manifestation of the theoretical requirement of "logical omniscience" implicitly assumed by S1 (see, Stalnaker 1991, Smithies 2015). Logical omniscience requires that P_0 satisfies the rules of probability calculus across the entire ■-algebra of propositions, $\mathcal{F}$. This requirement carries significant logical demands, which can be analytically broken down into three key aspects: (i) *Identification of Tautologies:* The agent must be able to identify all tautological propositions - those that are true in every possible scenario, such as "A or not A". Logical omniscience requires that the agent assigns a probability of 1 to every such tautology. (ii) *Detection of Contradictions*: Similarly, Logical omniscience dictates that the agent recognizes all contradictions - propositions that are false in every possible scenario, such as "A and not A" and assign probability zero to each one of them. (iii) *Recognition of Logical Entailments*: Beyond tautologies and contradictions, the agent must also correctly evaluate all logical entailment relationships between propositions. If proposition A logically entails proposition B, then $P_0(B)$ cannot be less than $P_0(B)$. Crucially, these theree requirements extend across the entire probabilistic space, demanding that the agent account for all possible logical relationships between propositions, not merely isolated subsets. This is the theoretical requirement of logical omniscience, mentioned above. Dutch Book arguments attempt to support this requiremnet by showing that even if an agent is not genuinely logically omniscient, her aversion to financial loss will compel her to form a system of beliefs equivalent to one shaped under logical omniscience. However, as we have shown, for this argument to succeed, the agent must instead possess computational omnipotence. Hence, the Dutch Book argument does not resolve the problem of the realism of S1; instead, it merely trades the theoretical requirement of logical omniscience for the practical requirement of computational omnipotence. While the former demands flawless logical reasoning across an entire probabilistic space, the latter requires the agent to possess unlimited computational capacity to detect and correct incoherences in any conceivable Dutch Book, no matter how complex. Both requirements seem to be far removed from the cognitive limitations of real-world agents.

3.2 Bayesian Conditionalization (S2): Epistemic and Pragmatic Arguments

The most important and convincing justification of Bayesian Conditionalization (S2) lies in the concept of *dynamic consistency*. An agent who updates her beliefs via Bayesian Conditionalization ends up with beliefs that are dynamically consistent. This means that the conditional belief formed at time $t = 0$ for the proposition A given the proposition E will be equal to her unconditional belief of A set at a later time t (where $t > 0$), provided that the event E is observed at t. This dynamic consistency ensures a coherence between the agent's current and future selves. Her initial expectations about how she would respond to evidence align perfectly with her actual response when that evidence is observed. But is this dynamic consistency a sign of rationality? The answer is yes, and it can be justified by two key arguments: one *epistemic*, proposed by Van Fraassen (1984, 1989), and one *pragmatic*, initially presented by Teller (1973) and later developed further by Lewis (1999).

Van Fraassen's epistemic argument for Conditionalization is rooted in the principle of *diachronic coherence*. According to van Fraassen, rational belief updating requires logical consistency across time. If an agent X assigns a conditional probability $P_0(A \mid E_t)$ at time $t = 0$, and later at time t learns that E_t is true, then her updated belief in A must match her earlier conditional probability. Failure to do so results in a form of epistemic inconsistency: X's past and present selves would hold contradictory beliefs about the same proposition, undermining the internal coherence of her belief system across time. For van Fraassen, this coherence is not merely a technical requirement but a fundamental norm of rational belief. The consistency between conditional probabilities assigned in the past and unconditional probabilities assigned in the future is what ensures that X's belief system remains epistemically sound over time. Conditionalization emerges as the *only* belief updating rule that can guarantee this kind of diachronic coherence.

Van Fraassen demonstrates that BC entails another principle of diachronic rationality, known as the *Reflection Principle (RP)*. RP may be roughly stated as follows: If X currently believes that at a future time t she will assign a probability of p to a proposition A and she trusts her future reasoning process, then her current probability of A should be p. RP encapsulates the idea that a rational agent should have consistency between her current and anticipated future beliefs, provided she has confidence in her future self's reasoning and ability to process evidence correctly. It ensures that there is no internal conflict between X's temporal belief states.

The pragmatic justification of Bayesian Conditionalization (S2) rests on the foundation provided by both the Dynamic Dutch Book Argument (Lewis, 1999) and the Converse Dynamic Dutch Book Argument (Skyrms, 1987). Lewis demonstrates that if an agent updates her degrees of belief in a way that violates BC, she becomes vulnerable to a series of bets spread across time that guarantee a loss, regardless of the outcome of events. Skyrms complements this by showing that if an agent updates her beliefs according to BC, she is immune to any diachronic Dutch Book. Together, these arguments establish an *if and only if* relationship: an agent avoids diachronic incoherence and vulnerability to Dutch Book exploitation if and only if she updates her beliefs according to BC.

However, the realism of this justification faces significant challenges. The computational burden of detecting complex betting strategies, already problematic in static Dutch Books, carries over to the dynamic case. Additionally, the assumption that an agent can foresee all possible evidential scenarios from time $t = 0$ and pre-assign conditional probabilities to them is problematic. Bacchus, Kyburg, and Thalos (1988) point out that BC assumes an idealized agent who can flawlessly anticipate every evidential scenario and pre-commit to precise conditional probabilities. In reality, agents are constrained by finite computational resources and are frequently unable to pre-encode conditional probabilities for every evidential possibility. Furthermore, belief revision in practice often involves a mix of updating and reflective re-evaluation, with agents revisiting their priors in light of new evidence, rather than merely applying a rigid conditionalization rule. This reflects a broader issue with the assumption that real-world agents can adhere to the strict rational requirements of Bayesian updating in all contexts.

3.3 Principal Principle (S3):

Deference to objective probabilities - the core prescription of PP - seems so intuitively compelling that it almost feels as though it needs no justification. Anyone who has consistently bet on games of chance understands this instinctively. To illustrate, consider agent X, who is about to enter a series of bets on a simple game: a coin flip. If the coin lands on heads, she wins one euro; if it lands on tails, she loses one euro. Crucially, the coin is fair, which means the objective chance of heads is 0.5 and the objective expected return of each flip is zero.

Now, imagine that X defers to the objective chance and adopts a subjective belief that mirrors the coin's objective probbailities. Her subjective expected return for each bet will also be zero. Over many flips, the actual return of the game will converge to zero due to the Law of Large Numbers. The agent's subjective expected return aligns perfectly with the actual return of the game in the long run.

Why is this alignment a good thing? Because it ensures that X's beliefs are *calibrated with reality*. Her expectations are neither overly optimistic nor unnecessarily pessimistic. If she had adopted a subjective belief that deviated from the objective chance - for example, if she believed the chance of heads was 0.7 instead of 0.5 - her expectation will systematically deviate from reality, leading her to make poor betting choices and potentially incur significant losses in the long run.

This example illustrates a *consequentialist justification* for PP: aligning subjective probabilities with objective chances works. It produces desirable outcomes, prevents systematic errors in reasoning, and ensures that subjective expectations are calibrated with the world's probabilistic structure. In this sense, adherence to PP is not presented as a principle of pure logic but rather as a pragmatically successful rule of rational action. If we follow PP, the argument goes, we will adopt subjective probabilities that motivate actions aligned with our long-term interests.

However, as Strevens (1999) observes, pragmatic success does not automatically confer logical necessity. The fact that adherence to PP yields good outcomes does not, on its own, constitute a foundational justification for why agents are rationally required to follow it. Just as the intuitive appeal of deference to objective chances cannot serve as a substitute for logical self-evidence, a consequentialist justification - rooted in the practical benefits of aligning belief with chance - remains incomplete if it cannot address the deeper normative question: Why should an agent defer to objective chances as a matter of rational necessity, rather than mere pragmatic convenience? Strevens (1999) argues, any attempt to provide a non-circular logical justification of PP ultimately fails. The core issue lies in the fact that such justifications inevitably presuppose the very principle they aim to establish. For example, consider the following argument (reliability-based justification): *Premise 1*: Objective chances reliably predict long-run frequencies. (Empirical Claim). *Premise 2*: If something reliably predicts long-run frequencies, then rational agents should align their subjective probabilities with it. (Normative Claim about Reliability ■ Rational Alignment). *Conclusion*: Therefore, rational agents should align their subjective probabilities with objective chances (PP). The circularity lies in Premise 2, where the transition from reliability to normative alignment implicitly assumes a version of PP or something very close to it. This circularity means that no purely logical argument can independently ground PP without smuggling in the principle's core assumption. Strevens concludes PP cannot be justified as a matter of logical necessity without falling into circular reasoning. However, the fact that we cannot non-circularly prove PP does not imply that it is not a real and robust standard for rational belief formation. Its pragmatic reliability and alignment with our best probabilistic practices ensure that PP remains a deeply rational principle, even if its justification ultimately rests on its success rather than a foundational logical proof.

A non-circular but metaphysically loaded argument for aligning our subjective probabilities with the corresponding objective ones is the following (ontic justification): *Premise 1*: It is rational to align one's beliefs with the true features of the world. (Normative Claim about Rationality ■ Alignment with Reality) *Premise 2*: Objective probabilities are true features of the world. (Metaphysical Claim – Ontic Interpretation of Objective Probabilities) *Conclusion*: Therefore, it is rational to align one's subjective probabilities with objective probabilities. (PP). This argument is not circular for the following reasons: Premise 1 is a general, independently plausible principle about rational belief formation: it's rational to align your beliefs with the truth. Premise 2 is a metaphysical claim asserting that objective probabilities are true features of the world. Hence, the conclusion follows from combining a normative principle about truth-tracking beliefs with a metaphysical fact about probabilities. Crucially, neither premise presupposes PP or any equivalent normative principle about aligning beliefs with chances. Therefore, this argument avoids circularity. However, it it relies heavily on the ontic interpretation of objective probabilities. If probabilities are indeed intrinsic features of an indeterministic world, the conclusion follows naturally as an imperative for rational belief. However, if objective probabilities are interpreted epistemically - as rational tools for managing uncertainty rather than as features of reality - the argument loses its force because epistemic probabilities are not "true features of the world" in the ontological sense required by the first premise. Thus, while the argument works well under an ontic interpretation, it collapses under an epistemic interpretation, highlighting that its success depends entirely on one's metaphysical stance on the nature of objective probabilities. This dependency reveals both the strength and the vulnerability of the justification: it is logically coherent but metaphysically contingent.

To sum up: while the ontic justification avoids circularity, it stands or falls with its metaphysical foundation. If one accepts that objective probabilities are true features of an indeterministic world, PP follows naturally as a principle of rational belief. On the other hand, the reliability-based justification, though pragmatically compelling, remains circular because it implicitly assumes what it sets out to prove. Despite these challenges, we believe that PP remains an indispensable principle of rationality - pragmatically robust, normatively compelling, and central to our best practices in reasoning under uncertainty.

3.4 Justifications of S4 and S5

The assumptions S4 and S5 represent two distinct but interconnected claims about how agents relate to the true probabilistic structure of the world. S4 asserts that the agent knows the true hypothesis H, which entails the true chance function $Ch(\cdot)$, from the outset. This assumption implies an idealized state of epistemic omniscience, where agents possess perfect and immediate access to the underlying probabilistic structure governing outcomes. In contrast, S5 adopts a more modest but still significant epistemic stance: the true hypothesis H lies within the set of alternative hypotheses $\{H_1, H_2, \ldots, H_n\}$ that the agent can conceive. Under S5, while the agent does not start with perfect knowledge of the true hypothesis, they have the conceptual resources to approximate it through observation, experience, and consistent updating.

The plausibility of S4, however, raises serious concerns. The assumption that agents possess perfect knowledge of the true hypothesis and the associated chance function is epistemically indefensible. Real-world agents are inherently bounded in their cognitive and informational capacities, making the assumption of perfect initial knowledge unrealistic. When advocates of rational expectations were pressed to justify this assumption, they frequently resorted to "arbitrage arguments." (see Muth 1961, Hausman 1989). Arbitrage arguments suggest that if agents' subjective probabilities systematically deviate from the true probabilities, opportunities for profit (or arbitrage) will arise. Rational agents, motivated by self-interest, will exploit these arbitrage opportunities, leading to adjustments in their probabilistic models. Over time, through a process of iterative correction and alignment with observable outcomes, agents' beliefs will converge toward the true probabilities, and only then will arbitrage opportunities be fully eliminated. Crucially, this justification hinges on a dynamic learning process rather than static omniscience, aligning far more naturally with the assumption embedded in S5 than with S4. Arbitrage arguments do not imply that agents begin with perfect knowledge; instead, they suggest that agents can, under certain conditions, refine their beliefs and approximate the true hypothesis over time. More specifically, arbitrage arguments operate under a pragmatic framework that does not require a pre-conceived hypothesis space. The agent does not explicitly formulate a set of alternative hypotheses. Instead, she reacts to observable mismatches between her probabilistic beliefs and real-world outcomes. If an agent's subjective probabilities systematically deviate from the true probabilities, this creates arbitrage opportunities (e.g., mispriced assets in financial markets). This implies that arbitrage arguments can be employed to provide an "as if" justification of S1, S2 S3 and S5 - but not S4. The term "as-if" is crucial here: agents do not need to explicitly satisfy these assumptions, nor do they need to reason within a hypothesis space or apply Bayesian updating consciously. Instead, their behavior - driven by profit opportunities and corrected by market feedback - produces outcomes that mimic the predictions of these assumptions. However, this dynamic adjustment process cannot replicate the epistemic claim embedded in S4, where agents are assumed to possess immediate and perfect knowledge of the true hypothesis from the outset."

Revisiting S4 in light of these considerations, it becomes apparent that its invention was likely driven by mathematical tractability rather than epistemic plausibility. In the canonical rational expectations models, terms such as $\mathcal{E}(y_{t+1} \mid I_t)$, denoting agent's subjective expectation of y_{t+1} based on the information available at t, posed significant mathematical challenges due to the subjective nature of expectations. The question was how should this term be treated mathematically? What are its properties? By equating agents' subjective expectations with the mathematical expectations $E(y_{t+1} \mid I_t)$ derived from the true model, economists were able to exploit the well-known properties of expectation operators, such as linearity and the law of iterated expectations, to achieve analytical tractability. This mathematical shortcut allowed for the closure of recursive systems and facilitated the derivation of equilibrium solutions. However, this methodological maneuver was subsequently dressed in the guise of an epistemic claim, creating the impression that agents possess an almost god-like knowledge of the probabilistic structure of the world. In reality, S4 should be viewed as a simplifying assumption - a heuristic device that served practical modeling purposes - rather than a plausible description of agents' epistemic capacities.

4 Conclusions

The purpose of this paper was to examine the alternative concepts of rationality employed in economics and to clarify their underlying epistemic assumptions and pragmatic justifications. We identified three main concepts of rationality: one based on coherence (Rationality-A), which emphasizes internal consistency of subjective probabilities for every point in time; one based on perfect accuracy (Rationality-B), which requires subjective probabilities to align exactly with objective probabilities at all points in time; and one based on asymptotic accuracy (Rationality-BB), where subjective probabilities converge to objective probabilities over time.

To support these concepts, we identified five key epistemic assumptions: S1, which ensures the coherence of the agent's initial probability function P_0 at $t = 0$; S2, which formalizes belief updating through Bayesian Conditionalization; S3, which mandates alignment with objective probabilities via the Principal Principle; S4, which assumes the agent knows the true probabilistic model from the outset; and S5, which posits that the agent operates within a hypothesis space that includes the true probabilistic model.

From these assumptions, we derived the logical relationships connecting them to the three concepts of rationality. Coherence for every t, captured by Rationality-A, follows from the assumptions of initial coherence and Bayesian updating (S1 and S2). Perfect accuracy, encapsulated by Rationality-B, requires not only coherence and updating but also adherence to the Principal Principle and knowledge of the true hypothesis (S1, S2, S3, and S4). In contrast, asymptotic accuracy, represented by Rationality-BB, relaxes the strong knowledge requirement of S4 and instead relies on the inclusion of the true hypothesis within a hypothesis space (S5), combined with Bayesian updating and alignment with objective probabilities (S1, S2, and S3).

Finally we examined the pragmatic justifications for these assumptions . These justifications aim to show that while real agents do not possess the extreme epistemic properties implied by S1–S5, they nevertheless - on pain of irrationality - act as if they had these epistemic properties. Arguments of this type typically fall into two broad categories: Dutch Book arguments, which aim to justify static and dynamic coherence (S1 and S2), and arbitrage arguments, which aim to justify asymptotic accuracy (S5). These justifications highlight how agents' behavior under certain constraints can mimic the outcomes predicted by these assumptions, even if they do not explicitly reason in the idealized ways implied by them.

However, the assumption S4 - that agents know the true hypothesis outright - is highly implausible. Real agents face cognitive and informational limitations, making perfect initial knowledge unattainable. The emergence of S4 in the rational expectations literature stems from its role in ensuring mathematical tractability. By equating subjective expectations with mathematical expectations derived from the true model, economists could exploit properties like linearity and the law of iterated expectations to solve recursive models. While this shortcut facilitated theoretical progress, it should not be mistaken for a realistic epistemic assumption.

In conclusion, while S1, S2, S3, and S5 find some pragmatic and epistemic support, S4 remains an idealized assumption reflecting methodological convenience rather than a defensible epistemic claim. This distinction underscores the importance of carefully separating theoretical assumptions from their pragmatic justifications when evaluating the foundations of rational expectations theory.

5 References

References

[1] Bacchus, F., H. Kyburg and M. Thalos (1988), "Against Conditionalization", Synthese, 85: 475-506.

[2] de Finetti, B. (1974), *Theory of Probability*, Vol. 1. New York: John Wiley and Sons.

[3] Hayek, A. (2008), "Arguments for – or against – Probabilism?.", *British Journal for the Philosophy of Science*, 59: 793-819.

[4] Kemeny, J. (1955), "Fair Bets and Inductive Probabilities", *Journal of Symbolic Logic*, 20: 263-273.

[5] Lehman, R. S. (1955), "On Confirmation and Rational Betting", *The Journal of Symbolic Logic*, Vol. 20(3): 251-262.

[6] Lewis, D. K. (1980), "A Subjectivist's Guide to Objective Chance" in his, *Philosophical Papers*, 2: 83-132, Oxford: Oxford University Press

[7] Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning.* Oxford University Press.

[8] Savage, L. (1954): *The Foundations of Statistics*. New York: John Wiley and Sons.

[9] Skyrms, B. (1987), "Dynamic Coherence and Probability Kinematics", *Philosophy of Science* 54: 1-20.

[10] Teller, P. (1973), "Conditionalization and Observation", *Synthese* 26: 218-258.

[11] Strevens, M. (1999), "Objective Probability as a Guide to the World", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 95: 243-275

[12] van Fraassen, B. (1984), "Belief and the Will", *Journal of Philosophy* 81: 235-256.

[13] van Fraassen, B. C. (1989), *Laws and symmetry*. Clarendon Press

6 Appendix

(1) Proof of Proposition 1:

The proof of this proposition is quite straightforward. We want to show that for every t, P_t obeys the rules of the probability calculus, namely non-negativity, normalization, and additivity.

(i) Non-negativity

From S1, $P_0(A \cap E_t) \geq 0$ and from S2, $P_0(E_t) > 0$. Therefore:

$$P_t(A) = \frac{P_0(A \cap E_t)}{P_0(E_t)} \geq 0.$$

(ii) Normalization

Consider the entire sample space Ω:

$$P_t(\Omega) = \frac{P_0(\Omega \cap E_t)}{P_0(E_t)} = \frac{P_0(E_t)}{P_0(E_t)} = 1.$$

(iii) Additivity

For two disjoint events A and B:

$$P_t(A \cup B) = \frac{P_0((A \cup B) \cap E_t)}{P_0(E_t)}. \tag{5}$$

Since A and B are disjoint, their intersections with E_t are also disjoint:

$$(A \cup B) \cap E_t = (A \cap E_t) \cup (B \cap E_t).$$

By additivity of P_0:

$$P_0((A \cup B) \cap E_t) = P_0(A \cap E_t) + P_0(B \cap E_t).$$

Substitute into (5):

$$P_t(A \cup B) = \frac{P_0(A \cap E_t)}{P_0(E_t)} + \frac{P_0(B \cap E_t)}{P_0(E_t)} = P_t(A) + P_t(B).$$

(2) Proof of Proposition 3:

Assume that at time t, the agent X observes the aXissible evidence data I_t. Given that $\mathcal{H}$ is a partition and (S2),

$$P_t(A) = P_0(A \mid I_t) = \sum_{i=1}^{n} P_0(A \mid H_i, I_t) P_0(H_i \mid I_t). \tag{6}$$

Admissibility of I_t implies,

$$P_0(A \mid H_i, I_t) = x_{A,i}. \tag{7}$$

Substituting $P_0(A \mid H_i, I_t)$ from (7) in (6) yields,

$$P_t(A) = P_0(A \mid I_t) = \sum_{i=1}^{n} x_{A,i} P_0(H_i \mid I_t). \tag{8}$$

Moreover,

$$P_0(H_i \mid I_t) = \frac{P_0(I_t \mid H_i) P_0(H_i)}{\sum_{j=1}^{n} P_0(I_t \mid H_j) P_0(H_j)}.$$

Substituting the last expression for $P_0(H_i \mid I_t)$ in (8),

$$P_t(A) = \sum_{i=1}^{n} x_{A,i} \frac{P_0(I_t \mid H_i) P_0(H_i)}{\sum_{j=1}^{n} P_0(I_t \mid H_j) P_0(H_j)} = \sum_{i=1}^{n} x_{A,i} \frac{1}{1 + \sum_{j \neq i} \frac{P_0(I_t \mid H_j) P_0(H_j)}{P_0(I_t \mid H_i) P_0(H_i)}} \tag{9}$$

The term $P_0(I_t \mid H_i)$ in (9) is the likelihood of I_t under H_i, $i = 1, 2, ..., n$. Given that X respects PP, and given that each H_i entails an objective probability $z_{I_t,i}$, respectively for I_t,

$$P_0(I_t \mid H_i) = Ch_i(I_t) = z_{I_t,i}. \tag{10}$$

Hence,

$$P_t(A) = \sum_{i=1}^{n} x_{A,i} \frac{1}{1 + \sum_{j \neq i} \frac{z_{I_t,j} P_0(H_j)}{z_{I_t,i} P_0(H_i)}}. \tag{11}$$

Ww are interested in the asymptotic behaviour of P_t, that is for $t \to \infty$. To that end, let us first assume that the experiment is repeated under independent and identically distributed (*i.i.d.*) conditions.[3] In this case, the law of large numbers applies, which means that relative frequencies converge to the corresponding objective probabilities. As a result, we obtain the so-called "likelihood ratio

[3] The "*i.i.d*" assumption is overly strict and is made only for simplicity. The results that follow can be extended to cover non *i.i.d.* cases.

convergence". Specifically, assume that H_1 is the true hypothesis and H_j, $j = 2, 3, ..., n$ are all false hypotheses and consider the following ratio,

$$\frac{z_{I_t,1}}{z_{I_t,j}} \times \frac{P_0(H_1)}{P_0(H_j)}. \tag{12}$$

Since H_1 is the true hypothesis, the ratio $\frac{z_{I_t,1}}{z_{I_t,j}} \to \infty$ (equivalently, $\frac{z_{I_t,j}}{z_{I_t,1}} \to 0$), as more information accumulates. This means that

$$\begin{aligned} \frac{1}{1+\sum_{j\neq i}\frac{z_{I_t,j}P_0(H_j)}{z_{I_t,i}P_0(H_i)}} &\longrightarrow 1,\ i = 1 \\ \frac{1}{1+\sum_{j\neq i}\frac{z_{I_t,j}P_0(H_j)}{z_{I_t,i}P_0(H_i)}} &\longrightarrow 0,\ i = 2, 3, ..., n, \end{aligned}$$

and finally from (11) and the above asymptotic relationships,

$$P_t(A) \to x_{A,1},\ A \in \mathcal{F}.$$

The last equation shows that, given that I_t describes a long series of outcomes, X's subjective probability of A is approximately equal to the objective probability, $x_{A,1}$, that the true hypothesis H_1 assigns to A. Moreover, X's credence in A does not depend on her priors $P_0(H_i)$. This result is usually referred to as the "washing-out of priors".