`

Arbitrage and Rational Decisions **by Robert Nau, Fuqua School of Business, Duke University**
Published by Chapman and Hall, 2025
https://people.duke.edu/~rnau/book/Arbitrage_and_Rational_Decisions.htm

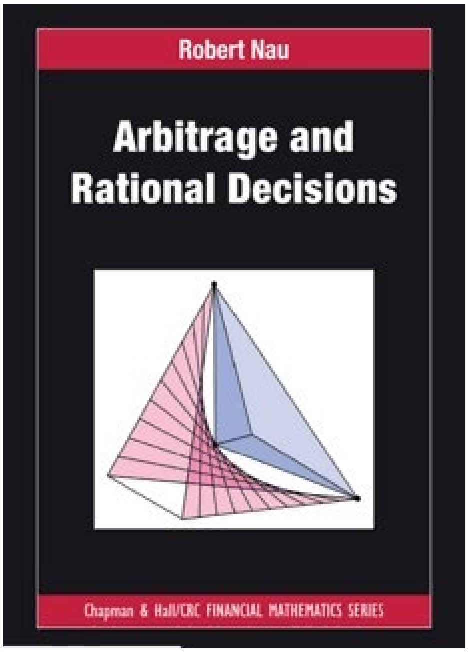

CHAPTER ABSTRACTS

Chapter 1 provides a survey of the book and it begins with the observation that rational choice theory—the theory of rational decisions by individuals and groups and societies---may be properly viewed as "social physics". It seeks universal mathematical laws that apply to social phenomena, and its key contributors throughout centuries have included great mathematicians and physicists of the day. The book provides a unifying perspective in which rationality concepts in different branches of the field are all characterized by a single principle, namely that decision makers should avoid exposing themselves to certain loss (arbitrage) at the hands of observers of their behavior. In each of the settings this principle has an equivalent ("dual") representation in terms of mental parameters such as probabilities, utilities, and strategic reasoning. This way of framing the theories highlights a number of fundamental issues. One is that it makes a great difference whether money is or is not involved as far as the use of precise numbers is concerned. Another is that, in general, observable economic behavior does not reveal subjective probabilities that can be interpreted as measures of pure belief, but (contrary to popular conception) this is not necessarily a problem. Yet another is that, at its most elementary level, noncooperative game theory can be founded on the same principle so that the existence of an equilibrium emerges as a result rather than a primitive requirement of rational strategic behavior. No-arbitrage is not only a standard of rationality: it is a standard of common knowledge of strategic rationality.

Chapter 2 begins with some history of developments in the theories of subjective probability and expected utility that occurred in the early to mid-20th century, following in the tradition of axiomatic methods of mathematics and physics. Systems of axioms were devised to model various properties of rational beliefs and preferences so that numerical representations of them could be derived by theorem-proving rather than direct assumption. The "independence" axiom, which yields representations that are additive in nature, was recognized to be fundamental and widely considered to be normative at that time, although it subsequently became controversial. (It holds that preferences among multi-faceted alternatives do not depend on elements that they have in common.) The chapter goes on to discuss the two most important tools in the rational choice toolkit: the fixed-point theorem and the separating-hyperplane theorem. The fixed-point theorem is the basis for proofs of existence of particular forms of equilibria in games and competitive economies. The separating-hyperplane theorem is the basis of proofs that observable behavior that does not lead to arbitrage (or which satisfies other kinds of consistency axioms) can be represented by probabilities and utilities inside the decision maker's mind. It is “where the numbers come from”, and in simple (finite) settings those numbers can be computed by applying linear programming models. The chapter concludes with a survey of fundamental representation theorems that span most of the spectrum of rational choice, emphasizing that they are all relatively minor variations on the single principle of no-arbitrage.

Chapter 3 deals with subjective probability theory: the theory of rational personal beliefs. The no-arbitrage standard of rationality is very familiar in this context: under the title “coherence” it is a cornerstone of the field of Bayesian statistics. A decision maker's subjective probabilities for events can be elicited either in terms of binary preferences among lottery tickets or in terms of acceptable bets. The latter approach has the advantage that the measurements involve a material interaction with an observer, which renders the results common knowledge between them and yields a material test for rationality: no-arbitrage. By the separating-hyperplane theorem, there is no arbitrage if and only if there is a subjective probability distribution by which lottery tickets and bets are evaluated according to their expected payoffs. Probabilities elicited in this way must comply with Bayes' rule, but only as a static condition, not a model of learning over time. A relaxation of the "completeness" axiom to allow different buying and selling prices for a lottery ticket leads to a more realistic representation of beliefs by intervals of probabilities rather than unique values. Incompleteness of preferences and non-uniqueness of mental parameters will be a recurring theme. An assignment of subjective probabilities avoids "ex-post-arbitrage" in a given state if it does not admit the acceptance of a bet with a negative payoff in that state and non-positive payoffs otherwise. By a variation on the main theorem this is true if and only that particular state has strictly positive subjective probability. This same result is the basis of the fundamental theorem of noncooperative games, which leads to Aumann’s concept of correlated equilibrium rather than to Nash equilibrium as a solution concept.

Chapter 4 deals with expected utility theory, where objects of choice are lotteries in which subjectively evaluated personal consequences are experienced with objectively given probabilities. This is the mirror image of subjective probability theory, whose objects of choice are lotteries in which objectively given consequences (amounts of money) are received with subjectively evaluated probablities. The two theories have same sets of axioms for their respective objects of choice, and these lead to similar additive representations of preferences: objective expected utility in the EU model and subjective expected value in the SP model. A behaviorally realistic relaxation of the completeness axiom leads to convex sets of utilities, rather than point values, analogous to what happens in the SP model. In the special case where the set of consequences is an interval of real-valued amounts of money, an infinite-dimensional version of the separating-hyperplane theorem can be used to derive the existence of a continuous utility function that models aversion to risk. Despite their mathematical similarities, the SP and EU models are operationally and philosophically very different. In the SP model, objects of choice are things that really exist in nature, their parameters can plausibly be common knowledge, and irrationality has material consequences. The generic EU model does not have these properties. Expected utility theory also provides a foundation for "social aggregation" theorems in which choices by a social planner are Pareto optimal (a type of no-arbitrage condition) if and only they are represented by a weighted sum of utility functions of individuals. The proof is (again) an application of the separating-hyperplane theorem.

`

Chapter 5 moves on to subjective expected utility theory, a synthesis of SP and EU theory in which both probabities of events and valuations of outcomes are subjective. A commonly-used way to set up the model is to let the objects of choice be state-dependent objective lotteries over a set of consequences. A straightforwad application of the EU axioms leads to a representation of preferences by a state-dependent utility function in which probabilities and utilities are not uniquely separated. An additional axiom is needed for that, along with a tacit counterfactual assumption that the same state of the person (consequence) can be experienced in the same way in any state of nature. The difficulty of separating beliefs about events from state-dependent tastes for consequences is emphasized throughout the book. In the state-dependent-utility version of the model, relaxation of the completeness axiom merely leads to convex sets of state-dependent utilities rather than unique values. Relaxation of completeness in the state-independent-utility version of the model leads to some difficulties. If no more than this is done, the representation of preferences may not consist merely of a set of state-independent utility functions (probability-utility pairs) or by separate sets of probability distributions and utility functions that can be paired up arbitrarily. Further constraints on preferences are needed to get those results. Also, as in the case of the EU model, the counterfactualism of the objects of choice and the personal nature of consequences raise questions in regard to whether one agent can know the parameters of another's preferences or that those parameters can furthermore be common knowledge, issues that are important in game theory.

Chapter 6 introduces state-preference theory, an alternative framework for modeling choice under uncertainty in which money plays an important role and objective probabilities do not play any role. Objects of choice are assignments of amounts of money (and perhaps other measurable attributes of outcomes) to states of the world. It does not aspire to the same generality as subjective expected utility theory, but neither does it invoke the contemplation of impossible private experiences. As such it provides a more robust basis for modeling interpersonal interactions and interpersonal knowledge. It has essentially the same axioms as subjective probability theory, minus the axiom of linear extrapolation of acceptable bets. Properties of preferences are measured by local betting rates for money, so-called "risk-neutral probabilities" (a term that is borrowed from financial asset pricing). Local attitudes toward uncertainty and ambiguity are measured by derivatives of the risk-neutral probabilities, which describe the curvature of indifference curves in all directions. If the independence axiom is not also imposed, the decision maker may display source-dependent attitudes toward uncertainty. A generalization of the Arrow-Pratt measure of risk aversion has a simple formula (a quadratic form) in terms of the matrix of derivatives of risk-neutral probabilities (a novel construct), and it does not presume expected-utility preferences or state-independent utility for money nor does it use imaginary riskless positions as reference points. The role and properties of risk-neutral probabilities and their derivatives are greatly underappreciated. Their applications to the modeling of ambiguity, noncooperative games, and asset pricing are discussed in the following three chapters.

Chapter 7 deals with aversion to ambiguity, that is, aversion to vagueness in probabilities, an issue raised in a famous paradox posed by Daniel Ellsberg. This has been a topic of huge interest in the field over the last several decades. According to subjective expected utility theory, uncertainties for all kinds of events can be represented by precise probabilities, and the decision maker is equally averse toward betting on any of them, regardless of their nature. Yet often the probability evaluations are based on differing quality of information and/or distorted by psychological or strategic considerations, and this can lead decision makers to be more averse to some sources of uncertainty than others. There are a great many ways to model this phenomenon, and the one that fits naturally into the state-preference modeling framework is to merely allow smooth indifference curves to bend in ways that violate the independence axiom. The

`

chapter shows that properties of the decision maker's attitudes toward uncertainty can be quantified in terms of measurable risk-neutral probabilities and their derivatives in a way that can explain phenomena such as Ellsberg's paradox . This method applies even if the decision maker has state-dependent utility for money and/or unknown prior stakes in events, issues that might otherwise be problematic. (The ability to address those issues is a hallmark of this approach.) Two different axiomatic models of smooth ambiguity-averse preferences are discussed, one that involves multiple sources of uncertainty, with different uncertainty attitudes applied to them, and another that involves second-order probabilities and utilities that model aversion to uncertainty about first-order probabilities as in the popular "KMM" model.

Chapter 8 deals with noncooperative game theory, the great middle ground of rational choice theory in which interlocking decision problems faced by small numbers of strategic competitors are modeled. It is the key chapter insofar as it shows how game theory can be directly unified with the theory of rational individual decisions and the theory of rational asset pricing in financial markets. Normally the modeling of rational game play begins with the introduction of additional principles of strategic rationality, such as an a priori notion of an equilibrium among the beliefs and intentions of the players. This chapter shows that when information about the rules of the game is revealed through incentive-compatible bets, the no-arbitrage principle leads immediately (via the ex-post version of the fundamental theorem of subjective probability) to the result that the apparent beliefs of the players must be consistent with a correlated equilibrium, a generalization of the solution concept of Nash equilibrium that is normally regarded as the bedrock of noncooperative game theory. Thus, the fundamental theorem of games is shown to be a small variation on the fundamental theorem of subjective probability and also the fundamental theorem of asset pricing. The chapter pursues various implications and applications of this idea, illustrated by the analysis of some famous examples, one of which--the so-called "battle of the sexes"--is the basis for the cover illustration on the book, which is shown above. (He prefers the boxing match, she prefers the ballet, but they would like to go somewhere together rather than separately. Where should they meet, and how, and why? The red saddle is the set of independently randomized strategies. The blue hexahedron is the set of correlated equilibria. Their three points of intersection—the black dots--are Nash equilibria. The obvious fair solution--flipping a coin--is the midpoint of the long edge, which is *not* a Nash equilibrium.) A new issue that emerges is that when players are risk averse (as modeled with diminishing marginal utility for money), the parameters of the equilibria become risk-neutral probabilities, which are the basis of models of asset pricing by arbitrage in financial markets.

Chapter 9 addresses the role of the no-arbitrage principle in the modeling of asset prices, a classic application area. According to the fundamental theorem of asset pricing (another multi-agent variation on the fundamental theorem of subjective probability), there are no arbitrage opportunities in the market if and only if there is a probability distribution according to which prices of assets such as stocks are the discounted expected expected values of their future payoffs. This probability distribution, the so-called risk-neutral distribution, is the common property of all the agents in the market. It behaves like the subjective probability distribution of a "risk-neutral representative investor", but it does not represent anyone's actual subjective probabilities, nor the average subjective probabilities of all agents, all of whom are risk averse. Rather, for each agent it is the product of her subjective probabilities and her state-dependent marginal utilities for money, where the latter are determined by her personal tolerance for risk and her personal portfolio of investments, which tend to cause systematic downward distortions in mean values. In a no-arbitrage state, the agents hold portfolios that equalize their risk-neutral probabilities with those of everyone else. There is fundamental indeterminacy in the true subjective probabilities of the real agents, as well as those of the average agent, which is the same problem that was encountered in subjective expected utility theory and state-preference theory. The effects of personal probabilities and utilities cannot be uniquely separated. The chapter goes on to consider the important

`

and especially tractable special case of normal probability distributions, exponential utility functions, and quadratic portfolios of investments, from which a purely subjective version of the capital asset pricing model (CAPM) can be derived.

Chapter 10 provides a summary of the key ideas and key models discussed in the preceding chapters, a largely non-technical helicopter tour. It is intended to serve both as a set of concluding remarks and as a condensed version of the book itself that will be accessible to a more general readership. Its first section provides an overview of foundational issues in rational choice theory, such as the definition of economic rationality (of which no-arbitrage is a unifying principle), the precision with which mental parameters such as probabiities and utilities can be measured, the difficulty in separating beliefs and tastes through observations of behavior, and the role of money in thinking and interacting in terms that have practical meaning and which can be quantified with precision, thus supporting the use of powerful mathematics. The chapter also emphasizes the limitations of the theory, especially the fact that it is only a low-dimensional static projection of a complex evolving social and economic world in which "rational" means more than "consistent." Section 2 shows the "dual" ways of modeling rational individual behavior (axioms of binary prefences and acceptable bets) that are the basis of subjective probability, expected utility, and subjective expected utility theory. The remaining sections walk through the contents of the other chapters, including the most important theorems, with a minimum of mathematical notation.

Chapter 11 presents the details of "linear programming" (constrained linear optimization) models that can be used to apply the fundamental theorems in the different application areas. The data for a given model consists of a matrix whose rows are payoff vectors of bets or trades that are acceptable to the agent(s). Each such model has a primal and a dual version. The primal version looks for a set of mental parameters (probabilities for events, utilities for consequences, state-dependent utilities for consequences, risk-neutral probabilities for payoffs, probabilities for game outcomes) that rationalizes the acceptable bets or trades by assigning them all non-negative value. The dual version looks for an arbitrage opportunity, i.e., a weighted sum of the acceptable bets or trades that is a sure loss for the individual or the group. In the solution of each model, only one of the two versions has a non-zero solution, i.e., either the right mental parameters exist or else there is an arbitrage opportunity. Models of this kind can be solved in practice by using the "Solver" tool in Excel.