



What are axiomatizations good for?

Itzhak Gilboa^{1,2} · Andrew Postlewaite³ · Larry Samuelson⁴ ·
David Schmeidler²

© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Do axiomatic derivations advance positive economics? If economists are interested in predicting how people behave, without a pretense to change individual decision making, how can they benefit from representation theorems, which are no more than equivalence results? We address these questions. We propose several ways in which representation results can be useful and discuss their implications for axiomatic decision theory.

Keywords Axioms · Axiomatization · Representation theorem · Decision theory

We are grateful to many colleagues for discussions that motivated and influenced this note. We thank Thierry Marchant, Ivan Moscati, Peter Wakker, and a referee for helpful comments and suggestions. ISF Grant 704/15 (Gilboa) and NSF Grants SES-0961540 (Postlewaite) and SES-1459158 (Samuelson) are gratefully acknowledged.

✉ Larry Samuelson
Larry.Samuelson@yale.edu

Itzhak Gilboa
tzachgilboa@gmail.com

Andrew Postlewaite
apostlew@econ.upenn.edu

David Schmeidler
davidschmeidler@gmail.com

¹ HEC, Paris, Jouy-en-Josas, France

² Tel-Aviv University, Tel Aviv, Israel

³ University of Pennsylvania, Philadelphia, USA

⁴ Department of Economics, Yale University, New Haven, USA

1 Introduction

Axiomatic decision theory was pioneered in the early twentieth century by Ramsey (1926) and de Finetti (1931, 1937) and achieved remarkable success in shaping economic theory. Bolstered by the axiomatic systems of von Neumann and Morgenstern (1944), Savage (1954), and Anscombe and Aumann (1963), expected utility became the dominant model of individual decision making in economics. When the critiques of Allais (1952, 1953) and Ellsberg (1961) blossomed into a concerted effort to generalize or develop alternatives to expected utility, axiomatic foundations again played a key role. A remarkable amount of economic research is now centered around axiomatic models of decision, both in the classical framework of choices between lotteries or between “Savage acts”, and in other models of risk and uncertainty.

What have these axiomatizations done for us lately? Are they leading to advances in economic analysis, or are they perhaps attracting some of the best minds in the field to deal with difficult problems that are of little import? Why is it the case that in other sciences, such as psychology, biology, and chemistry, such axiomatic work is so rarely found? Are we devoting too much time to axiomatic derivations at the expense of developing theories that fit the data?

This paper addresses these questions. Section 2 defines what is meant by an “axiomatization”—an axiomatization is an equivalence result relating a theoretical description of decision making to conditions on observable data. Section 3 cites some standard justifications of the axiomatic exercise and sharpens the questions of the previous paragraph to our single central concern: how can economists who study how people behave benefit from an equivalence result?

Section 4 provides our response, namely that axiomatic derivations are powerful rhetorical devices, and outlines several ways that axiomatic derivations of decision rules may be useful for economics, even when the decision models are interpreted descriptively.¹ Some real and imaginary case studies are discussed in Sect. 5, illustrating the role of axioms.

This discussion suggests two criteria, presented in Sect. 6, for judging which axiomatic work is most likely to be useful. Specifically, we hold that axiomatizations of general-purpose conceptual frameworks are probably more useful than axiomatizations of specific theories and that axiomatic models that describe rational behavior are likely to be more compelling than those describing specific instances of irrationality. Some concluding comments are offered in Sect. 7.

2 What are axiomatizations?

2.1 Representation theorems

In the context of decision theory, an “axiomatization” refers to a mathematical theorem that relates a decision rule involving theoretical concepts to conditions on presumably observable data. The best known and influential examples are the axiom-

¹ See Dekel and Lipman (2010) and Luce et al. (1990) for related discussions.

atization of the decision rules of maximization of utility (Debreu 1954; Jaffray 1975), of expected utility under risk (von Neumann and Morgenstern 1944; see also Herstein and Milnor 1953 and Marschak 1950), and of subjective expected utility under uncertainty (Savage 1954). In the first case, the derived theoretical concept is “utility” and in the second both “utility” and “subjective probability” are derived. In these examples, the presumably observed behavior is captured by a binary relation over alternatives, though other assumptions are also possible.

Such axiomatizations are also referred to as “representation theorems”. Typically, these also include results about the uniqueness of the mathematical representation of the data. In the famous examples cited above, the utility function for choice under certainty is unique only up to increasing transformations; the utility function for choice under risk or uncertainty, in the context of expected utility maximization, is unique up to positive affine transformations, and so forth.

There are many other usages of axiomatic systems. Some are less relevant to economics, such as axiomatizations of mathematical structures (the natural numbers, Euclidean geometry, etc.), and some belong to economic theory proper, such as axiomatizations of coalitional (cooperative) game theory or social choice solution concepts (the Shapley value, utilitarian social choice functions, and so on). These are not the topic of this paper. We focus here on the axiomatizations of decision models, relating theoretical constructs and choice procedures to presumably observable data.

2.2 What is presumably observable?

The assumption that a certain type of data “exist” in the sense of being observable is a primarily theoretical claim, made by scientists about reality. The reality consists of sources such as the Consumer Expenditure Survey published by the Bureau of Labor Statistics, reports of experiments that are run in labs, and so on. Scientists write a model in which the data are represented by formal objects, such as binary relations or choice functions. As in any act of modeling, mapping reality to formal objects typically leaves some room for interpretation. There can thus be a debate about what is observable. This is a scientific debate about the connection between data and models. It lies in the realm of the philosophy and methodology of economics more than in economics per se, but it is still scientific in that it deals with the degree to which formal models fit reality.

The question of the type of data that can be assumed observable in economics and decision theory has received much attention in recent years. For example, many textbooks in microeconomics present consumer theory as a special case of decision theory, where data are modeled as a binary relation, and this relation is then used to state conditions such as completeness, transitivity, and so forth. But two strands in consumer theory, motivated by considerations of observability, replace the notion of a binary relation by other formal models. One approach—“revealed preference”—argues that economists cannot observe choices between arbitrary bundles and that they can only observe selections from budget sets. In a

sense, this approach holds that a binary relation is assuming too much. Another approach—“stochastic choice”—holds that real data actually contain much more information than a binary relation admits, because choices are probabilistic (whether between two alternatives or, more generally, out of finite sets of alternatives). Thus, this approach suggests that binary relations assume too little. In both cases the formal model differs from the binary relation model in what it assumes observable.²

It should be taken for granted that the question of observability is a matter of degree, not of kind. Formal models invariably make some idealizations. For example, a formal model may assume that all data points are fully measured, while in reality some data are missing and need to be ignored or filled-in in a speculative way. Similarly, a model may ignore measurement errors, which are inevitable in reality. Beyond these problems, theory may require some idealizations. For example, the revealed preference approach in consumer theory assumes that consumer choices are observable for any income and price vector. This is undoubtedly closer to available data on consumer choice than is a binary relation defined over pairs of bundles. Yet, in reality, only a small subset of income and price values actually appear in the data, and some of the other values are unlikely to ever appear. Echoing an observation made by Marshall (1920, Book III, Chapter IV.8), we do not expect to observe household demand for the case in which the price of a car is \$10 and the price of gasoline is \$5,000 a gallon. Most people would even be hard-pressed to provide any subjective estimate of their demand under such conditions. Indeed, any workable formal model of the available data has to make some concessions to elegance and generality at the expense of realism.

Questions of observability are also intimately related to questions of relevance and external validity. For example, economists tend to mistrust non-choice data, such as self-report measures. Sometimes the issue at stake is the construct that is being measured (as in the case of subjective well-being), and sometimes it is the validity of the answers given in a questionnaire and their relevance to actual choices (as in the case of measuring the value of endangered species).

Economists tend to prefer axiomatizations that are based on data that are observable, but also that are as close as possible to the types of data economic analysis attempts to explain. For example, an axiomatization of subjective probability which is based on a self-reported “more likely than” relation over pairs of events is generally considered to be less satisfactory than another, based on a “preferred to” relation over pairs of acts. This preference may be justified by doubting the reliability of self-reported likelihood judgments (which can be viewed as a question of observability) or by questioning their relevance to real economic choices (which is akin to external validity). In any event, for the discussion that follows it is important to recall that an axiomatization makes not only explicit assumptions about regularities in the data, but also implicit ones about the data available.

² It is interesting to observe that both approaches date back to the 1950s at the latest (see, e.g., Luce and Suppes 1965 on stochastic choice), but that they recently seem to have moved to center stage after several decades in which they received less attention.

3 Goals of axiomatizations

What are axiomatizations good for? Relatedly, one may ask, why do we see so many representation theorems in the economic literature? We distinguish between three types of justifications. The first has to do with the methodology and philosophy of economics; the second with normative applications of decision theory; and the third with descriptive decision theory.

3.1 Meta-scientific

One reason to seek axiomatizations that relate theoretical concepts to observables goes back to the logical positivists' Received View (summarized in Carnap 1923) and to Popper's (1934) notion of falsifiability. Taken together, we obtain two dicta that were adopted by economists in the 1920–1950s: any theoretical concept should be related to observables, and any theory should be falsifiable. Much like the often-raised example of physics, and much unlike the severely criticized example of psychoanalysis (see Loewenstein 1992), the discipline of economics embraced the axiomatic approach. Axioms on observed choice ensured that theoretical concepts referring to mental phenomena that are not directly observable, such as “utility” and “probability”, have concrete meanings as well as algorithms for their measurement. This in turn guaranteed that theories formulated in these terms would have clear empirical implications and terms of falsifiability.³

3.2 Normative

When considering normative decision theory, as in the case of normative economics at large, the role of axioms is obvious and compelling. They can help the economist (or decision theorist) to convince the people she addresses that they would indeed like to follow her recommendation, or can call attention to weaknesses of a model. To consider a simple example, suggesting to people that they should maximize a utility function might seem weird in various contexts, especially when emotions are considered. Many listeners might suspect that the idea of quantifying anything and everything, and of maximizing functions when human values or emotions are at stake is the result of over-mathematization, showing poor understanding of human nature if not worse. Simple axioms such as completeness and transitivity, by contrast, are likely to be accepted by most listeners. A mathematical theorem stating that these two axioms are equivalent to behavior that can be described by utility maximization is then a powerful rhetorical device. It may at the very least present the utility maximization approach in a kinder light. Further, it can also convince some decision makers to

³ We refer to this type of justification as “meta-scientific” as it involves modeling of the work of scientists, as described in Sect. 2.2. above. It is not part of the work of the applied economist who attempts to explain or to shape economic reality; rather, it is part of the work of the philosopher or sociologist of science, who attempts to verify that the aforementioned economist does not use ill-defined terms and does not develop irrefutable, meaningless theories.

make the extra step and attempt to assess their own utility functions in the hope of guaranteeing that their choices are consistent.

Similarly, axiomatic systems such as von Neumann and Morgenstern (1944) or Savage (1954) might convince people that the way *they would like to make decisions* in the face of risk or of uncertainty is by maximization of expected utility. Clearly, not all axiomatic systems are equally compelling, and not all those that are compelling in some contexts are also in others.⁴ But there is little doubt that axioms are of paramount importance for normative applications of decision theory.

We mention in passing that the same argument would apply to axiomatic systems in social choice, welfare economics, coalitional game theory, intertemporal choice and related areas. In these fields the axioms are not used for representations of presumably available data by theoretical models; rather, they are about social choice and attempt to capture considerations of fairness, optimality, etc. Yet, the value of axioms as a rhetorical device in these cases is similar to the case of normative decision theory.

3.3 Descriptive

The question we wish to focus on is whether representation theorems are useful to economics as a descriptive science. After all, the meta-scientific interpretation of such results is methodological and of little relevance for most working economists. And, while the normative applications of decision theoretic models are sometimes important for economics (as in guiding government decisions), many economists view themselves as dealing only with positive questions. Further, even when normative economics is concerned, most economic analysis takes individual agents' preferences as given, as in the case of optimal taxation. Thus, for many if not most economists the normative justification of representation theorems has limited appeal. Economists would readily see why management, as an academic field, should be interested in such results, but would be left wondering why should they.

It is common to argue that axiomatic systems can prove useful also for descriptive purposes. Indeed, such systems can help delineate the scope of applicability of a given theory and can serve as guides for the development of experiments to test it. When a theory is refuted by data, axioms may help in identifying which parts of the theory are the weaker ones and might, therefore, be the first to be replaced or generalized.

Observe, however, that all of these goals are actually meta-scientific. They do not say something about the reality modeled, but about the work of the scientist who models it. A common approach to assessing a theory is to break it into its more manageable axiomatic components and then to investigate these components. But the axioms themselves have no effect on the content of the theory or how well the theory serves as a description. We can thus finally sharpen our question: do representation theorems help economics as a descriptive science, and if so, how?

It would seem that there is a simple argument for a negative answer: representation theorems are mathematical characterization results. Ideally, they show the equivalence of two mathematical representations of the data. For example, any data set of choices in

⁴ Specifically, we refer to Gilboa et al. (2009, 2012) who discuss the appeal of Savage's axioms and the importance of the state space to which they are applied.

which completeness and transitivity hold can be represented by utility maximization, and vice versa. Thus, whether a given data set refutes the theory of utility maximization is independent of the way this theory is represented. If we have panel data, with multiple choices made by different individuals, if we find that 78% of the individuals do not refute utility maximization, we know that 78% of them do not refute the axioms that are equivalent to utility maximization. If we wish to compare the theory of utility maximization to a competing theory (say, Simon's satisficing, Simon 1957), the theory that fits the data better would not depend on its representation. Indeed, a decision theorist who attempts to improve upon a theory might be interested in the axioms it satisfies; but why would an economist who only wishes to use the theory for economic predictions or recommendations be interested?

4 Rhetoric

Casual observation, as well as introspection, suggests that axiomatizations enhance and deepen our understanding of economic models. They provide insights into these models and help economists realize what their models actually assume. Admittedly, characterization theorems are useless for those who can immediately see all tautologies, but the rest of us can use an axiomatization in order to better understand "what they buy" when they adopt a specific model.

The question still remains, however, how do understanding and insights help economists in better describing an economic environment. With or without enhanced understanding, a theory is just as accurate a description of the data in any mathematically equivalent representation. How do the insights provided by an axiomatization result in better predictions, then?

We argue that axiomatizations, and the insights they provide, can be powerful rhetorical devices. We use here "rhetoric" not in the negative sense, referring to tricks that are designed to win an argument, but in the positive sense, referring to tools of reasoning that a person would be truly convinced by, and might use herself next time she debates the issue.⁵ The most powerful tool one can use to convince others is logic, or, more generally, mathematics. Axiomatizations, relying on mathematical theorems, can potentially be useful rhetoric.⁶ While the value of rhetoric is clear for normative applications, it is not so obvious for descriptive ones. Our question is, therefore, why do we let axiomatizations convince us that a theory is more or less plausible? Why not circumvent the axioms and compare the theory directly to the data? We describe three scenarios in which axiomatization may be useful in such a debate, alongside available data.

⁵ This is in line with our views of economic models as rhetorical devices (see Gilboa et al. 2014, 2106). Our focus here, however, is on the choice of modeling tools rather than on specific models.

⁶ Axiomatizations that are effective as rhetorical devices can also be thought of "framing effects": a decision rule that may appear unreasonable in one representation may be more compelling in another.

4.1 Undeveloped theories

One case in which scientists may select models by their inner logic or elegant properties rather than by their fit to the data and parsimony occurs when theories have not yet been developed, and the scientists are looking for a conceptual framework within which theories will be developed. In this case we think of axiomatic systems as a way for one economist to convince herself or others that a particular model is more promising than another. The question, “why not see which theory fits the data better” is answered, quite simply, by “it’s too early to tell”.

4.2 Failed theories

It is widely accepted that concrete, quantitative theories in the social sciences cannot boast the level of accuracy that is attained by theories in the natural sciences. In some cases there is an acceptance of the fact that all such theories are bound to have counterexamples. What does this leave economists with? In particular, when asked to make predictions about markets or growth, or to provide advice about interest rates and unemployment benefits, what theory of individual behavior can economists employ to base their answers on? In such circumstances, economists are led to ask, which of all the wrong theories is perhaps the least wrong for the purpose of the discussion at hand? And in making such judgments, axiomatic systems may prove helpful.

4.3 Untestable theories

A related scenario in which axiomatizations can be helpful arises out of the inability to directly test theories. For example, suppose that one is interested in traders’ behavior under stress, when the stakes are high. Suppose that empirical data do not allow an identification of specific factors such as risk aversion parameters. One may use experimental results in an attempt to isolate such a factor, but the external validity of the experiment would be low: the stakes and psychological stress would necessarily be much lower in the lab than in real life. Attempting to run a more realistic experiment, perhaps with a poorer population (for which the experiment’s stakes might be high) might run into ethical problems. In short, there is no practical way to test the desired effects in a lab. Here, again, axiomatic systems might be of help as a rhetorical substitute for data, helping economists convince each other that a certain theory makes (or does not make) sense.

Notice that the last two types of reasons often overlap. Specific theories of human behavior can often be tested, perhaps in lab experiments, and some of them can be reasonably accurate within their intended domain. However, economists are often interested in theories that have implications for situations that are inherently difficult to experiment with (such as macroeconomic events) and for phenomena in which individuals’ behavior is hard to isolate in a controlled way. These theories can be regarded as failed theories, because they falter in the situations to which they are applied, or as “untestable theories”, because these situations are so complicated as to render testing impossible.

5 Case studies

We devote this section to case studies, illustrating the three types of justifications described above. Each case is presented as a representative of one of the justifications, though it can also be viewed as illustrating others.

5.1 Expected utility theory: undeveloped theories

The debate during 1948–1951 over the theory of expected utility maximization under risk is a fascinating example of the role of axioms in choosing a conceptual framework before actual theories are developed within the framework. As Moscati (2016) describes, luminaries of economics, including Paul Samuelson, Jacob Marschak, and Harry Markowitz, did not see expected utility theory as *the* normative or descriptive model of decision under risk. They tended to view expected utility theory as an instance of convex preferences to which no special status should be attributed. They were convinced that expected utility should be the benchmark model by arguments, especially letters and personal communications from Leonard Savage, involving the von Neumann and Morgenstern (1944) axioms, including the ability to derive the independence axiom from Savage’s sure-thing principle.⁷

Imagine we could be privy to these discussions. The people involved were basically restructuring economic analysis, using mathematics and statistics in ways that had not been attempted before them. They surely knew that, in the best of cases, they were launching a long-term project that many young economists would join in generations to come. Thus, they were probably asking themselves, which tools should we equip the young generation with to cope with problems of economics? Constrained optimization and utility maximization were accepted as unrivaled tools for modeling rational behavior—but how does one deal with risk? Clearly, it was way too early at this point to say, “let’s see how this formula copes with problems of information economics, and to what extent it is a good approximation of human behavior”.⁸ Neither the explanatory power of expected utility theory for economic modeling nor its violations in experiments could be fully foreseen at the time. And, in the absence of data about the successes of the theories that would one day be developed using the conceptual framework, one had to find other ways of evaluating this framework. Axiomatic analysis proved to be a very powerful tool in this debate.

⁷ As pointed out to us by Ivan Moscati, these and other leading economists varied in their motivation and in their views of expected utility as a normative and/or descriptive model of choice. We thank a referee for pointing out that economists at the time were apt to blur the distinction between normative and descriptive models. For example, Arrow (1951, p. 406) writes that, “In view of the general tradition of economics, which tends to regard rational behavior as a first approximation to actual, I feel justified in lumping the two classes of theory together.” One contribution of Kahneman and Tversky (1979) was to argue that descriptive models can be improved by not insisting they be normative.

⁸ Moscati (2018) describes early empirical investigations of the expected-utility axioms.

5.2 Utility maximization: failed theories

Consider the following (imaginary) scenario: In the context of a debate over markets, an economist raises the first welfare theorem and points out that *competitive* markets have the advantage of leading to Pareto efficient allocations, where the latter are defined in terms of the agents' utility functions. There are many reasons why this argument may fail to convince a skeptic of the virtues of competitive markets. The skeptic may question the assumptions of the competitive model, such as price-taking behavior and perfect information. Similarly, the skeptic may question the assumption that each agent acts as if maximizing a utility function. Most people do not appear to constantly calculate utilities or take derivatives. Why should we think of them as utility maximizers? And if they do not maximize utility, we cannot even define Pareto optimality, let alone rely on agents making "optimal" choices given budget constraints. Then why should we heed the arguments of an economist who studies a market populated by utility maximizing agents?

The economist may parry this criticism of utility maximization with an appeal to axioms, arguing, "OK, admittedly, there are many assumptions involved in my argument. But the assumption that people behave as if they maximize a utility function isn't as demanding as it appears to be—I'm essentially arguing only that people's choices satisfy transitivity." And indeed, the claim that (for the most part and especially when important decisions are involved) people make transitive choices sounds much more compelling than that they maximize a utility function.

Observe that utility maximization has not proven to be a generally correct theory. If it had, the economist's argument using the first welfare theorem would have been much stronger. And, importantly, it would not have depended on the way the theory was represented, by a mathematical formula or by behavioral axioms. But this is not the case. There are too many violations of the basic assumptions of utility maximization to say that it has been "proven to be correct".⁹ Unfortunately, no other theory has proven to be correct either. The state of the art seems to be that any specific theory of behavior may at best "work" in some set-ups, but fail in others. Indeed, Amos Tversky used to say, "show me the axiom and I'll design the experiment that refutes it".

Given this state of the art, economists who are asked to make predictions and/or recommendations have to choose among theories that are known not to be as precise as one would have wished. And, while theories of individual decision making are typically tested in lab experiments, economists are asked to project from experimental results to relevant economic set-ups, using their judgment and common sense. When these subjective inputs are concerned, there is no surprise that axiomatizations can have an effect.

⁹ Among the best known violations of expected utility theory, preference reversals strike at the heart of the transitivity axiom. See Lichtenstein and Slovic (1971), Lindman (1971) and Grether and Plott (1979).

5.3 Prospect theory: untestable theories

An interesting example, illustrating both the need for axioms and its absence, is the development of prospect theory.¹⁰ Kahneman and Tversky (1979) proposed a concrete theory that was rather focused: it dealt with one-stage lotteries, each involving up to two values apart from the reference point, with known and clearly stated probabilities.

While Kahneman and Tversky offered a sketch of an axiomatization in their (*Econometrica* 1979) paper, it was probably not the most convincing argument in favor of the model. In particular, the theory presented in that paper assumed several stages—coding, combination, segregation, and cancellation—that were implicit in the model and were not axiomatized. Further, the axiomatic derivation focused on the case of “regular” prospects, in their equation (1), and did not go further to obtain their equation (2). Despite the limitations of the axiomatization, prospect theory has been very successful. The Kahneman and Tversky paper is one of most widely cited papers among those published in an economic journal over the past 50 years.¹¹

There clearly are many reasons for the success of this particular paper. Some of these may be more sociological: for example, it appears that this paper came to epitomize the general critique of models of rational choice theory. Yet, one reason for its popularity is rather classical: prospect theory is simply a successful theory. In the restricted set-up for which it was proposed, prospect theory was, and is considered to this day to be a very good approximation to the way people make risky decisions. Importantly, when the theory works, it can succeed without a compelling (or even complete) axiomatization.

However, for economic applications prospect theory needed to be extended to deal with more than three outcomes per prospect. How should that be done? One extension, which was considered by many to be what Kahneman and Tversky had in mind, was to follow Edwards’s (1954) formula. Edwards suggested that a lottery $(x_1, p_1; \dots; x_n, p_n)$ (guaranteeing outcome x_i with probability p_i) is evaluated by the decision maker by a function

$$V(x_1, p_1; \dots; x_n, p_n) = \sum_{i=1}^n f(p_i) v(x_i) \quad (1)$$

for a value function v (monotone in its real-valued argument) and a monotone function f , mapping objective probabilities into decision weights. One can then combine this idea with the notion of a reference point, considering, as does prospect theory, the x_i ’s to be changes in wealth relative to the reference point, rather than absolute values.

However, this formulation is not very promising. As Fishburn (1978) showed (commenting on a model of Handa’s (Handa 1977)¹²), this formulation violates first-order stochastic dominance unless f is linear, in which case expected utility theory re-emerges. By contrast, rank-dependent utility theory (Quiggin 1982; Yaari 1987) suggested that the function f be applied to decumulative probabilities so that the

¹⁰ This is a real case that has not been analyzed by professional historians of science. It is based on the authors’ own impressions and personal histories.

¹¹ We thank Peter Wakker for this and related observations.

¹² Handa (1977) suggested the same formula where v was the identity.

functional to be maximized is

$$V(x_1, p_1; \dots; x_n, p_n) = \sum_{i=1}^n \left[f\left(\sum_{j=1}^i p_j\right) - f\left(\sum_{j=1}^{i-1} p_j\right) \right] v(x_i), \quad (2)$$

where $x_1 \geq \dots \geq x_n$.¹³ Interestingly, the formulation of prospect theory for prospects with two positive or two negative outcomes was in line with (2) rather than with (1), with the latter used explicitly only for prospects with one positive and one negative outcome.

There is reason to believe that Kahneman and Tversky did not have (2) in mind for the general case, and later they adopted the rank-dependent approach. Combining it with the notion of a reference point, they suggested cumulative prospect theory (Tversky and Kahneman 1992). In addition, Tversky and Kahneman (1992) provide an axiomatization of cumulative prospect theory, opening the door to extensive testing of the theory. This axiomatization played a key role in the success of prospect theory.

Arguably, this history can be seen as an instance of axiomatizations helping out when theory cannot be directly tested. Kahneman and Tversky's preference for restricting attention to lotteries with small support was based on psychological considerations, preempting the possible critique that expected utility theory was violated simply because the choices presented to the participants were too complex. With some freedom of interpretation one could argue that, based on such considerations, any specific formula would be refuted in experiments involving complex prospects. That is, the general case of decision making between prospects with arbitrarily large supports can be seen as an instance of failed or of untestable theories. In a sense, one would be on safer ground by not venturing into the messy lands of complex stimuli. But if economic analysis requires a model to work with, it might be safer to generalize the experimentally tested model based on theoretical arguments rather than on experimental results that become too noisy. Thus, the axiom of first-order stochastic dominance provided a good argument in favor of cumulative prospect theory over prospect theory.

6 Which models should we axiomatize first?

The three reasons for which one may be interested in representation theorems for descriptive purposes may help us understand why in some cases the axiomatic exercise seems more compelling than in others. We believe that economists tend to see the rationale for such an exercise more readily when the model axiomatized is more general, and when it represents more "rational" behavior. We now turn to explain these terms more concretely and relate them to the discussion above.

While our main goal is to capture the reaction of fellow economists to axiomatic exercises, this section admittedly has a more normative tone than the previous ones. It should be clarified that any mathematically correct, nontrivial and/or surprising axiomatization theorem is of some value and may prove useful to economic analysis in one

¹³ This turned out to be a special case of Choquet expected utility theory (Schmeidler 1989).

way or another. The question we wish to address is not whether some axiomatizations are of any value, but which ones are likely to have a greater impact on economic analysis at large.

6.1 Generality

By generality of a model one typically refers to some measure of the set of its possible applications. The comparison of models by a “more general than” relation can thus be captured by set inclusion.

While this notion of generality appears to be a good approximation of what people have in mind when they think about scientists explaining data sets, it seems to be missing some important factors when the discourse among theorists is concerned. Consider the following examples:

Example 1 von Neumann and Morgenstern’s independence axiom (in the standard formulation, due to Jensen 1967; see also Fishburn 1978) states that

$$P \succ Q \Rightarrow \alpha P + (1 - \alpha) R \succ \alpha Q + (1 - \alpha) R$$

for all lotteries P, Q, R with finite support (and every $\alpha \in (0, 1)$). Assume that we apply this theory to outcomes that are integer monetary values (in a given currency). Imagine that independence is now weakened to P-Independence, which states that the above implication holds in case P, Q, R have supports contained in the set of prime monetary values, but not necessarily otherwise. Clearly, P-Independence is logically implied by Independence. The model that assumes that \succ satisfies P-Independence generalizes that which assumes that \succ satisfies Independence. And yet, this generalization raises suspicion.

Example 2 Decision theory deals with a single decision maker. Consider next a theory of games played among 23 players. Clearly, the latter generalizes the former. Any decision problem with a single decision maker can be viewed as a special case of a 23-player game, of whom 22 are “dummy” players. Yet, a theory of games played by 23 players would raise eyebrows, whereas decision theory will not.

These examples are obviously extreme. But they still raise the question: why do these generalizations appear so ludicrous? Why do not we feel the urge to work on such models, or even to read them?

The two examples are not of the same type. In Example 1 we discuss a more general theory, stated in a given language, whereas in Example 2 we discuss a more general model, using a richer language. Yet, we argue that there is a common thread to both.

Let us begin with Example 1. Clearly, the degree to which we believe in an axiom is not a monotone function of its logical strength. Part of the reason is that, to be convinced of an axiom, one has to go through some reasoning. If the axiom says, “Condition C holds”, one has to look for the logic of Condition C. If, by contrast, the axiom says, “Condition C holds under Circumstances D”, one first has to be convinced of the logic of Condition C, then consider some counter-examples, and then see that they are ruled

out by Circumstances D. That is, the mental process by which one imagines a set of instances might be longer than the process needed to imagine a superset thereof. This could not occur if scientists were thinking of all possible instances of a model, imagining them one by one, and considering how likely is a certain axiom in each of them. If we were thinking of sets of instances in this intensional way, a more general condition could only be considered to be less probable than a less general one. But imagining all possible instances is hardly a good model of how people think about axioms. Rather, they think in an extensional way, where the sets of instances are represented in some concise way. For that reason, conjuring up a set of instances may require a longer process of imagination than would a superset thereof.

Relatedly, reading the P-Independence axiom, a reader might say, “Why on earth would you restrict the independence axiom to lotteries with supports in the prime numbers? With all due respect for number theory, can anyone seriously expect an axiom to work for prime-valued lotteries and not for others? What made you come up with this weakening of the axiom in the first place?” This suggests that part of our faith in axioms has to do with our judgment of the motives of the person presenting them and the information available to them. We realize that scientists participate in some game, in which they try to construct new theories and get credit for them. In a way that parallels Grice’s Principle in the philosophy of language (Grice 1957), we expect scientists to come up with the simplest theories that work. Hence, when someone suggests a theory that is restricted to Circumstances D, we tend to make the inference that without this restriction the theory probably does not work, even if this was not explicitly stated. And then, when judging “Condition C holds under Circumstances D” we tend to ask ourselves, “OK, now that we infer that C does not hold universally, do we believe that assuming D is enough to rule out the counterexamples?” In some circumstances, if we have good reasons to explain why C would hold only under D, we may accept the weaker axiom. For example, assume that B-Independence is the von Neumann and Morgenstern Independence axiom limited to lotteries in which there are no probabilities in the range $(0, 0.1)$. Like P-Independence, this axiom also generalizes Independence, by limiting the set of instances that might count as a refutation. However, in this case we may come up with a reasonable explanation for the generalization: we have learned from experiments that people have problems understanding small probabilities; B-Independence tries to do away with the problematic small probabilities but retain the basic logic of the Independence axiom. There is thus some logic for the restriction of the axiom (and the generalization of the theory). By contrast, if, as in Example 1, the restriction of the axiom to circumstances D appears arbitrary at best, we have reason to suspect that there would be many other counter-examples to the general principle, and no good argument to believe that the current weakening of the axiom would indeed fare any better than its original version.

Let us now turn to Example 2. Here, again, the choice of the number of players in a game to be 23 appears arbitrary. In particular, 23-player games do not seem to be a conceptual framework, or a “paradigm”, which can be used for a variety of economic problem. Ironically, if we have to use such a framework to deal with an unknown problem, it seems safer to forget about 22 of the players, consider a single-person decision problem, and remind ourselves that games can be analyzed as special cases of decision problems, where the uncertainty involves the behavior of others. In other

words, a theory of a single decision maker can give readers a sense of a general conceptual framework, while a theory of 23-player games does not. One might not think decision theory is the ideal framework for analyzing games—perhaps because one thinks uncertainty about Nature should be modeled differently than uncertainty about optimizing opponents—but one may still prefer decision theory to the 23-player game. While the latter is a logical generalization of the former, it does not conjure in our minds a general method for modeling various situations of choice.

As in Example 1, we can also consider here the social aspect of scientific activity. A theorist who comes up with a model for a single decision maker seems to be honestly interested in a certain class of phenomena. By contrast, another who suggests a theory of 23-player games raises questions of motivation: why 23? Why not more, or less? Has the person tried the same principles with other games and failed miserably? Should we infer from what is said also about what is not?

We believe that most economists (the authors included) tend to prioritize axiomatizations of more general models over more specific ones. Considering the reasons given in Sect. 4 for which axiomatizations might be useful as rhetorical devices, the first one obviously favors generality: if the model axiomatized is to be considered a conceptual framework within which yet-unconceived-of theories are to be developed, its generality is of paramount importance. Also, if we have to choose a model where each specific theory is either untestable or refuted, a more general one would certainly have greater appeal. However, in all of the above we find that the notion of generality does not correspond to simply the set of possible applications. Rather, psychological and social considerations, relating to the mental process of imagination and to the scientific game, have to be brought to bear in judging generality of models.

6.2 Rationality

Faced with some of the axiomatic decision models in recent decades, more than one person apparently came up with the Tolstoy paraphrase, “All rational people are rational in the same way, but all irrational ones are irrational in their own way”. We mean by this that looking for axioms, which are general, abstract principles governing human behavior, seems more fruitful when the axioms describe behavior that can be thought of as “rational” than when it is “irrational”.

To make this argument meaningful, one has to define “rationality”. For many economists rationality is defined by some classical set of axioms, say, Savage’s (1954), and any behavior that veers away from the model is automatically dubbed irrational. This is not the way we use rationality here. Rather, we adopt a more subjective definition of the term. As in Gilboa and Schmeidler (2001), suppose that “rationality” refers to a robustness of sorts: a mode of behavior is rational for a decision maker if the latter does not wish to change it once exposed to its analysis. One reason the decision maker may not wish to change her behavior might be that she knows she cannot. For example, according to this definition it may be rational not to play chess optimally, or to fail to solve NP-Hard problems optimally. The decision maker also may not wish to change her behavior because she may not be convinced by theoretical arguments. For example, some people may prefer not to behave in accordance with Savage’s model

when it comes to choices involving global warming, because, despite understanding Savage's axioms, they do not feel comfortable with a seemingly arbitrary choice of a prior.

Gilboa (2009) and Gilboa et al. (2009) refined this notion of rationality and distinguished between objective rationality, which has to do with the ability to convince others that one is right, and subjective rationality, whose essence is the ability to withstand criticism by others and not be convinced that one is wrong. Both notions are based on the act of convincing others, and, at least in this respect, they bring to mind Habermas's (1981) notion of communicative rationality.

We have stressed the importance of axioms as rhetorical devices. They are supposed to be *convincing*. As such, they seem to be intimately related to the notion of rationality. Observe that the argument for rationality of axioms is much more powerful in the case of normative applications than of descriptive ones. Indeed, for a normative application the claim that axioms should be rational is almost a tautology: normative models are supposed to convince people that they would like to make decisions in a certain way, and if they can perform this task, they are objectively rational according to the definition above.¹⁴ This relationship is not as tight for descriptive applications. In these cases, it is not the decision maker who has to be convinced that they would like to make decisions in a certain way, but an economist who has to be convinced that *other* people tend to make decisions in a certain way. It is certainly logically coherent to believe, for example, that most people (including oneself) are prone to certain logical mistakes but not to accept these as desired goals.

At the same time, simple, elegant conditions that attempt to capture what most people would consider to be mistakes do not seem to be as compelling as similarly elegant conditions that are accepted as rational. For example, economists may readily accept *modus ponens* as a tenet of rationality, even while recognizing that all humans may violate it in their reasoning from time to time. However, someone interested in a model of irrational behavior would not typically accept an axiom stating that whenever a person accepts p and $p \rightarrow q$, they also accept $\neg q$, even if the axiom is softened by recognizing that it may be violated from time to time. Instead, one would like to study under which conditions *modus ponens* is violated. The answer would probably lie in detailed psychological studies, taking into account the person's age and education, the topic over which they reason, the time allowed and the stakes involved—in short, many factors that are typically absent from axiomatic decision theory.

As the notion of rationality we refer to is subjective, our argument cannot possibly classify axiomatizations as good or bad, worthwhile or worthless. However, we believe that the general reaction of economists to an axiomatic model is justifiably more positive the more rational is the behavior under discussion.

¹⁴ Clearly, "convincing people" should be operationalized in a quantitative manner. Some people may be convinced to follow a decision model, while others may not; they may accept it for some applications but not for others, etc. Note that these measurement issues apply both to the term "normatively appealing" and to "objectively rational".

7 Concluding comments

7.1 Interpretation in scientific analysis

In the axiomatizations we discuss, presumably observable choice behavior is shown to be equivalent to a representation that involves mathematical constructs such as “utility function”, “probability measure”, and the like.¹⁵ These constructs are named so as to suggest interpretations based on mental concepts such as “desirability”, “belief”, and so forth. In Economics, such an exercise is often taken to be a *definition* of the mental concept in question.

There are two types of critiques that can be raised against such a definition. First, it can be argued that, in practice, the axioms do not hold for many decision makers, or perhaps need not hold even for rational decision makers, and, therefore, the implicitly suggested mathematical measurement of the mental concept is groundless. Second, it is possible that the axioms hold but that the interpretation is not as compelling as it might seem. In the first case it is sometimes possible to re-define elements in the model so as to allow the axioms to hold, and, correspondingly, to revisit the interpretation of the mental concepts in a richer model. We illustrate this using two examples.

7.1.1 Example 1: utility under certainty

As is well-known, a complete and transitive binary relation over a finite set can be represented by maximization of a function, typically referred to as the “utility function”. Note that the name chosen and the way that the function is used in much of economic theory (including the welfare theorems) suggests that a higher utility is a good thing. In particular, obtaining a higher utility values for the functions of some individuals without decreasing those of others is considered desirable.

A possible critique of the first type was raised by Sen (1993), who argued that a binary relation is not always well-defined. In a classical example, he suggested to consider a decision maker who likes apples but knows she is not supposed to appear greedy, and, therefore, always selects the second largest apple from a choice set. Violating independence of irrelevant alternatives, such a decision maker cannot be described by a binary relation, or by a utility function that is defined on alternatives alone.

This critique can be countered by re-defining the objects of choice, so that not only apples, but also image and social standing are at stake. By including these relevant factors into the definition of “an outcome”, and allowing the utility function to depend on them, one can re-instate the binary relation and the interpretation of the utility function as measuring desirability. Observe that such a re-definition may be restricted to theoretical work, whereas in empirical work, in the absence of data on “image” and “social standing”, one may have to use a more general conceptual framework than simple utility maximization.

By contrast, a possible critique of the second type would be along the lines of the well-being and happiness literature, arguing that people may obey the axioms of utility

¹⁵ Other concepts can be “qualitative probability”, “cost of temptation”, “similarity”, etc.

maximization, but that it would be premature to interpret the mathematical function that describes choice as a measure of well-being. Specifically, people may prefer jobs that pay higher salaries, but the pursuit of riches does not lead to happiness. This claim is not about the descriptive or normative validity of the axioms, but about the meaning one should attach to the mathematical constructs used to describe behavior.

7.1.2 Example 2: subjective probabilities

Savage's derivation of subjective probabilities was criticized for assuming state-independent utilities. Critiques of the first type suggested that some of the axioms would not hold (descriptively) and even need not hold (normatively). Specifically, Aumann argues that Savage's axioms P3 and P4 are both likely to be violated if one considers outcomes whose desirability depends on the state of the world, such as a swimsuit and an umbrella.

Notice that, as in the case above, one can re-define the concepts to salvage the axioms. Indeed, Savage's reply to Aumann was that a swimsuit and an umbrella are not final outcomes, but uncertain acts, and an outcome should be sufficiently informative to determine the decision maker's well-being, as "lying by the pool on a sunny day", or "running for cover half naked in the rain".¹⁶ As above, such re-definitions of outcomes can be very useful for theoretical work, but may be of little help for empirical work relying on the scanner data about the purchase of swimsuits and umbrellas.

However, Savage's definition of subjective probability and state-independent utility was also subject to a critique of the second type. In another famous example by Aumann (see Karni 1996), a decision maker may satisfy all of Savage's axioms, but, because in one event (of his wife's death) his utility is a linear transformation of his utility in the complement event, the subjective probability "measured" by Savage's representation would not correspond to his actual beliefs. That is, the axioms might hold, giving rise to a unique probability measure-cum-state-independent utility, but the interpretation of the former as a representation of beliefs would be unwarranted.

7.1.3 Conclusion

One should be careful in interpreting the mathematical constructs that are derived from an axiomatic approach to behavior. The axiomatic project we discuss can sometimes be viewed as an attempt to measure presumably unobservable mental processes based on presumably observable behavior. Thus, choice provides meaning to the vague concept of "utility", betting behavior provides meaning to "subjective probability", and so forth. If the axioms fail to hold, one receives an alarm signal. If a decision maker violated the independence of irrelevant alternatives (under certainty) or Savage's P3 and P4 (under uncertainty), one can tell that there is a problem with the data. One approach would be to re-define concepts, and thus take into account important factors that affect the mental process under consideration. But if the axioms do hold, one

¹⁶ Famously, Savage also ridiculed his own defense by writing "I don't mind being hanged as long as my reputation and good health are unharmed". See Gilboa (2009) for further discussion of the delineation of applications in which state-dependent utility is unavoidable.

should be careful not to jump to the conclusion that the mathematical constructs derived from them necessarily have the intended meaning.

Note that in the above we use the prefix “presumably-” both to observed behavior and to unobserved mental processes. While economic theory sometimes errs in assuming that choice is more readily observable than it really is, it also sometimes errs in pretending that choice is all that is observable. In particular, self-report data are often useful, and they would prove relevant in both of our examples.¹⁷

7.2 Interpretation in meta-scientific analysis

Meta-scientific applications of axiomatizations also rely on rhetoric to a large extent. Consider, for example, the von Neumann and Morgenstern axiomatization of expected utility maximization under risk. It can be viewed as a definition of the theoretical concept “utility”, providing a way of measuring it, and a specification of the degree of uniqueness of the function.¹⁸ But in order to measure the utility function of a given individual one need not resort to the axioms; one may directly use the representation by the expected utility formula, as long as it is not refuted.¹⁹ However, the theorem can help convince economists that the term “utility” will not be devoid of content for many agents: whenever an agent satisfies the axioms, that agent’s utility will be measurable and well-defined. Such a claim can affect the economist’s faith in the meaningfulness of the term without necessarily testing it.

References

- Allais, M. (1952). Fondements d’une théorie positive des choix comportant un risque et critique des postulats et axiomes de l’École Américaine. *Econométrie*, 257–331.
- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’École Américaine. *Econometrica*, 21(4), 503–56.
- Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematics and Statistics*, 34, 199–205.
- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica*, 19, 404–437.
- Carnap, R. (1923). Über die Aufgabe der Physik und die Anwendung des Grundsatzes der Einfachheit. *Kant-Studien*, 28, 90–107.
- Cozic, M., & Hill, B. (2015). Representation theorems and the semantics of decision-theoretic concepts. *Journal of Economic Methodology*, 22, 292–311.
- Debreu, G. (1954). Representation of a preference ordering by a numerical function. In R. M. Thrall, C. H. Coombs, & R. L. Davis (Eds.), *Decision processes* (pp. 159–165). New York: Wiley.
- de Finetti, F. (1931). Sul significato Soggettivo della Probabilità. *Fundamenta Mathematicae*, 17, 298–329.
- de Finetti, F. (1937). La prevision: ses lois logiques, ses sources subjectives. *Annales de l’Institut Henri Poincaré*, 7, 1–68.

¹⁷ In the first, self-report has famously been used to measure well-being. Despite many problems with this measurement tool, “Subjective Well-Being” is still considered an important source of information. Likewise, in Aumann’s example involving the decision maker’s wife self-report would be a way to find out that the measured probability does not correspond to his beliefs.

¹⁸ See Moscati (2018) for a discussion of the development of utility theory.

¹⁹ See Cozic and Hill (2015) on axiomatizations as definitions of theoretical terms.

- Dekel, E., & Lipman, B. L. (2010). How (not) to do decision theory. *Annual Review of Economics*, 2, 257–282.
- Edwards, W. (1954). The theory of decision-making. *Psychological Bulletin*, 51, 380–417.
- Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Fishburn, P. C. (1978). On Handa's new theory of cardinal utility and the maximization of expected return. *Journal of Political Economy*, 86, 321–324.
- Gilboa, I. (2009). *A theory of decision under uncertainty*. Cambridge: Cambridge University Press.
- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy*, 25, 285–296.
- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2012). Rationality of belief. *Synthese*, 187, 11–31.
- Gilboa, I., Postlewaite, A., Samuelson, L., & Schmeidler, D. (2014). Economic models as analogies. *Economic Journal*, 128, F513–F533.
- Gilboa, I., Postlewaite, A., Samuelson, L., & Schmeidler, D. (2016). Economics: Between prediction and criticism. *International Economic Review* (2016) (forthcoming)
- Gilboa, I., & Schmeidler, D. (2001). *A theory of case-based decisions*. Cambridge: Cambridge University Press.
- Grether, D. M., & Plott, C. R. (1979). Economic theory of choice and the preference reversal phenomenon. *American Economic Review*, 69, 623–638.
- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66, 377–388.
- Habermas, J. (1981). *Theory of communicative action, Vol I: Reason and the rationalization of society, English edition* (p. 1984). Boston: Beacon.
- Handa, J. (1977). Risk, probabilities, and a new theory of cardinal utility. *Journal of Political Economy*, 85, 97–122.
- Herstein, I. N., & Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, 21, 291–297.
- Jaffray, J.-Y. (1975). Existence of a continuous utility function: An elementary proof. *Econometrica*, 43, 981–983.
- Jensen, N.-E. (1967). An introduction to Bernoullian utility theory I, II. *Sedish Journal of Economics*, 69(163–183), 229–247.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Karni, E. (1996). Probabilities and beliefs. *Journal of Risk and Uncertainty*, 13, 249–262.
- Lichtenstein, S., & Slovic, P. (1971). Reversals of preferences between bids and choices in gambling decisions. *Journal of Experimental Psychology*, 89, 46–55.
- Lindman, H. R. (1971). Inconsistent preferences among gambles. *Journal of Experimental Psychology*, 89, 390–397.
- Loewenstein, G. (1992). The fall and rise of psychological explanation in the economics of intertemporal choice. In G. Loewenstein & J. Elster (Eds.), *Choice over time* (pp. 3–34). New York: Russell Sage.
- Luce, R. D., Krantz, D., Suppes, P., & Tversky, A. (1990). *Foundations of measurement, Vol. III: Representation, axiomatization, and invariance*. New York: Academic Press.
- Luce, R. D., & Suppes, P. (1965). Preference, utility and subjective probability. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. III, pp. 249–410). New York: Wiley.
- Marschak, J. (1950). Rational behavior, uncertain prospects, and measurable utility. *Econometrica*, 18, 111–141.
- Marshall, A. (1920). *Principles of economics* (8th ed.). London: Macmillan.
- Moscati, I. (2016). How economists came to accept expected utility theory: The case of Samuelson and Savage. *Journal of Economic Perspectives*, 30, 219–236.
- Moscati, I. (2018). *Measuring utility: From the marginal revolution to behavioral economics*. Oxford: Oxford University Press.
- Popper, K. R. (1934). *Logik der Forschunghe: The logic of scientific discovery (English Edition 1958)*. London: Hutchinson and Co.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3, 225–243.
- Ramsey, F. P. (1926). Truth and probability. In R. Braithwaite (Ed.), *The foundation of mathematics and other logical essays* (p. 1931). London: Routledge.
- Savage, L. J. (1954). *New York: The foundations of statistics*. Oxford: Wiley.

- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57, 571–587.
- Sen, A. (1993). Internal consistency of choice. *Econometrica*, 61, 495–521.
- Simon, H. (1957). *Models of man*. New York: Wiley.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton: Princeton University Press.
- Yaari, M. (1987). The dual theory of choice under risk. *Econometrica*, 55, 95–115.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.