



Formal Epistemology Meets Mechanism Design

Jürgen Landes¹

Accepted: 5 August 2022 / Published online: 22 February 2023
© The Author(s) 2023

Abstract

This article connects recent work in formal epistemology to work in economics and computer science. Analysing the Dutch Book Arguments, Epistemic Utility Theory and Objective Bayesian Epistemology we discover that formal epistemologists employ the same argument structure as economists and computer scientists. Since similar approaches often have similar problems and have shared solutions, opportunities for cross-fertilisation abound.

Keywords Formal epistemology · Economics · Computer science · Mechanism design · Rationality · Game theory

1 Introduction

Formal epistemologists, much like non-formal epistemologists, are interested in an agent's knowledge and its rational and/or justified beliefs. Much of formal epistemology has been devoted to the idea that a rational agent's beliefs come in degrees. The most popular explanation of gradual beliefs are formal models of beliefs employing probabilities. Arguments to justify these models have been devised applying Dutch Book Arguments or scoring rules.

Economists and computer scientists have been much interested in properties of the allocation of goods and services to economic agents. Much work has been devoted to the notions of an efficient market and profit maximisation. A key insight is that allocations of goods and services do not appear magically but are part of a *market that can be designed*. In order to construct a market conducive to having desirable properties, economists and computer scientists created incentive structures for economic agents; a process referred to as *mechanism design*.

This article draws parallels between formal epistemology and mechanism design. Analysing key arguments in both fields, we discover the hitherto unobserved fact that researchers in both fields employ the same argument structure. This is a priori somewhat surprising, given that both fields are isolated from each other as evidenced by an absence of cross-referencing. Due to the unexpected closeness philosophical thinking can have immediate intellectual and practical impact. Furthermore, by engaging with the relevant literatures

✉ Jürgen Landes
juergen_landes@yahoo.de

¹ Department of Philosophy “Piero Martinetti”, University of Milan, Milan, Italy

philosophers can discover concepts, techniques and considerations which are useful to their own philosophical endeavours. On the other hand, researchers working on mechanism design may apply their advanced mechanisms fruitfully in formal epistemology.

Close connections and potential cross-fertilisation between formal epistemology and other branches of computer science (machine learning and artificial intelligence) have previously been identified (Ortner and Leitgeb 2011; Williamson 2004).

The rest of the paper is organised as follows. The next two sections are used to introduce mechanism design and formal epistemology (Sects. 2 and 3). Section 4 explicitly traces parallels between mechanism design and formal epistemology. Section 5 offers some conclusions.

2 Mechanism Design

The basic underlying idea of mechanism design is familiar to all of us: to bring about preferable circumstances. Next, we show how this basic idea has given rise to the field of mechanism design. By discussing one truly remarkable mechanism we will delineate an argument structure which is key in mechanism designing and which we re-discover in three areas of formal epistemology.

2.1 Idea, Challenge and the Internet

Economists have a long-standing interest in analysing markets in all their facets (equilibria, actors, collapses, cycles, etc.). In the 1990s, the idea took hold that markets do not simply just exist, markets can also be *created*. Hence, the creator of a market can influence market behaviour of market actors by the way the market operates. It is not uncommon for a market creator to also be a market actor, such creators hence have strong preferences about how the market operates. Consider for example an auction house, which designs its own specific rules to further its own interests. The challenge arises to create a market—or design a market mechanism—that satisfies (as much as possible) the designer’s preferences.

The main difficulty the designer faces is that—within the set boundaries—actors are free to act any way *they* choose. That is, once the market has been designed, the designer does not have means to *force* actors to behave a certain way. For example, the auction house cannot force potential bidders to submit (high) bids.

Ceteris paribus, economic actors prefer lower coordination costs. Information technologies, such as the internet, have long been recognised to have enormous potential for reducing coordination costs, see for example Malone et al. (1987). Unsurprisingly, computer scientists hence also became involved in mechanism design. Nowadays, much effort is expended on the design of tailor-made web-based market mechanisms.

Two successful applications of designed web-based mechanisms are reported in Hohner et al. (2003) and Sandholm (2007). The former reports how tailor-made auctions at Mars Incorporated benefited Mars *and* its suppliers. The second tale is an impressive story of selling an eye-popping \$35 billion through specially designed auctions. Mechanism design done well can be really beneficial to the designer and others. However, not all designed mechanisms are successful, the fitting titles of Cramton (2003), Klemperer (2002) are: “Electricity market design: the good, the bad, and the ugly” and “How (not) to run auctions”.

Overviews of scholarly work on mechanism design can be found in Cramton et al. (2006), Klemperer (1999), Nisan (2014), Roth (2002), Dash et al. (2003). The former two chart auctions while the latter two are more concerned with computational issues. A recent area of exploration are computational issues in auction design which is surveyed in Chawla and Sivan (2014), Hartline (2012).

We now visit a truly remarkable landmark mechanism, the Vickrey Auction.

2.2 Vickrey Auction

Let us consider a salesman wishing to sell one indivisible unit of good or one service, e.g., a painting. Clearly, he has a preference to sell at a high price but there are a number of other desires concerning a transaction mechanism that he may have. He might prefer to do business:

- (i) quickly (time is money),
- (ii) simply (so that everyone clearly understands how the transaction mechanism works),
- (iii) treat all potential trading partners equally (fairness),
- (iv) ensure that no actor revealing true preferences loses utility (keep future potential trading partners happy and honest),
- (v) allocate the good/service to the actor with highest private valuation—assuming all actors are rational (most beneficial to the economy as a whole, according to many economic theories),
- (vi) the outcome of the transaction mechanism is consistent with the private information of every actor (to qualm worries about fair play) and
- (vii) keep the private valuations of market actors private.

The Vickrey Auction, introduced in Vickrey (1961), satisfies all these desiderata. A Vickrey Auction is an ingenious single-round game for N players interested in buying one item from the salesman:

- (A) *Rules* Every player privately transmits one single bid (some positive number in \mathbb{R}) to the salesman.
- (B) *Outcome* The item is awarded to the player with the highest bid who publicly pays the *second* highest bid submitted.
- (C) *Utility* The utility of the highest bidder is the difference between her utility of the item and the price. The utility of all other players is zero.
- (D) *Decision norm* Players ought to bid such that their bid is not weakly dominated by another bid: avoid bids for which there exists another bid which never returns less utility and sometimes more utility.
- (E) *Mathematical theorem* The only way for a player to avoid a weakly dominated bid is to bid her private true valuation of the item. If all players bid true valuations, then i–vii obtain.

Let us briefly discuss why (i)–(vii) obtain. The game is

- (i) single-round (only one bid),
- (ii) it does not involve higher maths,
- (iii) its outcome only depends on the amount agents bid,

- (iv) neglecting the time and effort spent for putting in a bid: the utility of actors submitting unsuccessful bids is zero, the utility of the highest bidder is strictly greater than zero, it is the difference between her bid (equalling her true valuation) and the second highest bid (demonstrated below),
- (v) holds (demonstrated below) and
- (vi) also holds.
- (vii) What becomes public knowledge is that the agent allocated the item put in the highest bid and that the announced price was the second highest bid, nothing more.

For the following discussion we assume that no two players submit the exact same bid. Couldn't a player do better by bidding less than her true valuation in case she won the auction? If her new lower bid is not the highest bid any more, she is not allocated the item and her utility drops to zero. If she is still the highest bidder, then she still has to pay the same price (the second highest bid), in which case her utility remains unchanged. So, putting in a lower bid is not a good idea. What about a player bidding more than her true valuation, in case she was not the highest bidder? If the new higher bid is still below the maximal bid, then her utility is still zero. If the new higher bid becomes the highest bid, then the player has to pay a price which exceeds her valuation of the item which is not ideal. All other cases are similar.

Summing up, bidding true valuation is guaranteed to be as good as every other bidding strategy. Any other bidding strategy yields lower utilities in certain cases and it is never strictly better. This holds independently of what all the other bidding actors do.

But what about the salesman and his preference to sell at a high price? The item is traded for the second highest bid which equals, if all players are rational, the second highest true valuation. The item is hence sold for a high price if and only if at least two agents value it highly. Noting that any price greater than the highest true valuation is a salesman's pipe-dream, the second highest true valuation appears like one reasonable way to reflect the salesman's negotiation power. Clearly, if the salesman knows that there is only one actor who highly values the good or service, then it might make more economic sense to directly negotiate with this actor.¹

A lot of work has gone into modifying the Vickrey auction. For example, the celebrated Vickrey-Clarke-Groves (VCG) auction (Clarke 1971; Groves 1973; Vickrey 1961) generalises the above game to simultaneously selling multiple items; see Abrache et al. (2007) for an overview. This generalisation also satisfies (i), (ii), (iii), (iv), (vi) and (vii). As for (v): the distribution of the multiple items maximises the sum of private valuations.

Other directions of work are auctions designed for an agent seeking to buy—rather than sell—from a group of salesmen. Such auctions are called 'procurement' or 'reverse' auctions. One issue in such auctions is that items available may differ in quality and price simultaneously, refer to Branco (1997), Che (1993) for dated overviews and see Buettner and Landes (2012) for an online reverse VCG auction with differing qualities and prices, multiple bidders and multiple sellers as well as bids consisting of multiple components.

¹ Another option is to add a "reserve price" to the auction. If this price is not achieved, then the item will not be sold.

2.3 Structure of Mechanism Design

Economic actors (e.g., the salesman, procurer) have different preferences on how they want to do business. They hence want to use a mechanism tailored to their own preferences.

To achieve a particular set of economic goals in a particular economic situation, P , a mechanism designer follows a road map: design a game G specifying

- (A) rules,
- (B) outcomes,
- (C) utility and a
- (D) decision norm characterising rational action for agents playing game G .
- (E) Mathematical theorem: if all agents playing game G are rational in the sense of (iv), then P obtains.

We already saw that ensuring (i)–(vii), a *designer's* preferences in a particular economic situation, through the design of the Vickrey auction was along this road map.

2.4 Mechanism Design—Up Close and Personal

At this point, the philosophical reader might wonder what all of this has to do with *yourself*. Well, mobile-phone licenses are sold by governments through auctions which are particularly tailored to the occasion. Every time you make a call on a mobile you use a license bought through an auction (Klemperer 2002).²

Then there is eBay, a rather successful application of mechanism design. Reputation systems on eBay designed to take the function reputation plays in real-world markets are effective as reported in Houser and Wooders (2006), Resnick et al. (2006).

If the reader is at this point not persuaded that mechanism design is worth knowing about, let me briefly tell you that some have thought about how (not) to sell the monstrosities known as ‘nuclear weapons’ by designing a clever mechanism. The brave reader who wants to venture on such dangerous grounds is referred to Jehiel et al. (1996).

With this rough guide of the lands of mechanism design in mind, we now turn to philosophical shores.

3 Formal Epistemology

3.1 Some Background

Formal epistemologists, much like non-formal epistemologists, are interested in an agent's knowledge and its rational and/or justified beliefs. Much of formal epistemology has been devoted to the idea that a rational agent's beliefs come in degrees: a proposition F can be *believed to a degree*. The leading account of degrees of belief (often called *credences*),

² Early on, revenues from the auctions were mind-boggling to many. The, at the time, German minister of finance, Hans Eichel, re-named UMTS—originally abbreviating ‘Universal Mobile Telecommunications System’—to “*Unerwartete Mehreinnahmen zur Tilgung von Staatsschulden*”—which roughly translates to *Unexpected More-receipts to repay State debts*, see [Link](#).

Bayesianism, has become a paradigm in the philosophy of science (Easwaran 2011a; 2011b; Sprenger and Hartmann 2019; Landes 2021b; Joyce 2011; Huber 2016; Weisberg 2015; Radzvilas et al. 2022; Skipper and Bjerring 2022; Schupbach 2022).

One core tenet of Bayesian epistemology is *Probabilism*, the norm that credences ought to satisfy the axioms of probability and all such credences are, in principle, admissible. There is considerable debate as to which other—if any—further constraints apply to a rational agent's degrees of belief.

Another popular thought among formal epistemologists is that credences should not contradict *known chances* (sometimes referred to as *objective probabilities*). For example, if it is known that a coin is fair, then it seems unwise to believe that the coin to be tossed shortly will show 'heads' to a degree $\frac{1}{10}$. This idea underlies the *Principal Principle* of Lewis (1980).

Many writers have also held that a symmetry in an agent's body of evidence ought to entail a corresponding symmetry in credence. There is a number of formalisations of this principle, called the *Principle of Indifference*, on the market. In its weakest form this principle says: 'Over a given and fixed finite set Ω of N mutually exclusive and jointly exhaustive set of atomic propositions, a rational agent's credence in every $\omega \in \Omega$ ought to be $\frac{1}{N}$, if the agent does not possess any evidence concerning the $\omega \in \Omega$.' Arguments in favour of this principle have recently been offered in Novack (2010), Paris (2014), Pettigrew (2016b), White (2010), Eva (2019), Decock et al. (2016). Surprisingly, some Bayesians controversially argue that the Principal Principle implies the Principle of Indifference (Hawthorne et al. 2017; Landes et al. 2021b), others have disagreed (Pettigrew 2020; Titelbaum and Hart 2020; Gyenis and Wroński 2017).

Undoubtedly, Probabilism is more widely accepted than all rival formalisations of graded beliefs. Nevertheless, there are alternative ways to capture the idea that rational belief comes in degrees. One objection to Probabilism is that it is sometimes impossible to precisely specify degrees of belief in a proposition given the available evidence. In particular, the Principle of Indifference cannot be accepted. Rather, the agent's ignorance about an atomic proposition $\omega \in \Omega$ ought to be reflected by not assigning any precise definite credence to it. One such alternative is Dempster-Shafer-Theory, see Shafer (2011) for a modern treatment, and imprecise probabilities more generally, see Bradley (2015) for an overview. Alternatively, ranking functions have been suggested as a model for rationality (Spohn 2012).

Summing up, there are a number of (mutually conflicting) epistemic attitudes a rational agent may have. A philosopher may think that (in a particular epistemic situation) a rational agent possesses some specific subset of epistemic attitudes. It is incumbent on the philosopher to give an argument for her thoughts.

Unbeknownst to them, formal epistemologists employ the exact same argument structure used by designers of mechanisms in economics and computer science I outlined in Sect. 2.3. Formal epistemologists design mechanisms to satisfy their epistemological goals. An unexpected consequence of our analysis is that these philosophical arguments share the same structure.

3.2 Dutch Book Arguments

The, arguably, most famous argument in rational belief formation is de Finetti's *Dutch Book Argument* for Probabilism (de Finetti 1937; 1980; Pettigrew 2020b). De Finetti

designed a single-round, two-player, zero-sum game between a bettor and a book-maker as follows:³

- (A) *Rules* On a finite set of possible worlds, Ω , containing the actual world $\omega^* \in \Omega$, the bettor has to assign a betting rate to all $F \subseteq \Omega$, $b_F \in \mathbb{R}$. The book-maker chooses a stake $s_F \in \{-1, 0, +1\}$ for every $F \subseteq \Omega$; selling or buying the bet at the betting rate b_F .
- (B) *Outcome* At the world $\omega \in \Omega$, the bet on F returns $r_\omega(b_F) = s_F \cdot b_F$, if $\omega \notin F$ and $r_\omega(b_F) = s_F \cdot (1 - b_F)$, if $\omega \in F$.
- (C) *Utility* The bettor's utility at world ω is the return of all bets, $\sum_{F \subseteq \Omega} r_\omega(b_F)$. The book-maker's utility is $-\sum_{F \subseteq \Omega} r_\omega(b_F)$.
- (D) *Decision norm* The bettor ought to avoid sure loss, i.e., do not adopt credences b such that there exist stakes s_F for all $F \subseteq \Omega$ such that for all $\omega \in \Omega$ it holds that $\sum_{F \subseteq \Omega} r_\omega(b_F) < 0$.
- (E) *Mathematical theorem* All probabilistic and only probabilistic credences avoid sure loss. Furthermore, no probabilistic credences are loss dominated: for all different probabilistic credences b, b' there exist different $\omega, \omega' \in \Omega$ such that $\sum_{F \subseteq \Omega} r_\omega(b_F) > \sum_{F \subseteq \Omega} r_\omega(b'_F)$ and $\sum_{F \subseteq \Omega} r_{\omega'}(b_F) < \sum_{F \subseteq \Omega} r_{\omega'}(b'_F)$.

By altering the argument structure a philosopher can bring about that different sets of credences appear to be rational. For example, the Dutch Book Argument has been modified such that the set of 'optimal' credence functions are the intuitionistic probability functions, see Weatherson (2003), or the Dempster-Shafer-Belief functions, see Shafer and Vovk (2001). Worth pointing out is that de Finetti and Weatherson designed a mechanism for two players while Shafer and Vovk devised a three-player game.

A further variant of the Dutch Book is the Czech Book in which the decision norm is altered to seek bets which 'guarantee financial gain'. In this set-up, the purported rational credences are all those which are non-probabilistic, see Hájek (2008, 796-797). Extensions of the Dutch Book to groups of agents have been considered, see Kopec (2017) for some of the latest words. Dutch Book Arguments for different underlying logics were studied in Paris (2001).

There are two further credences which can be argued for by modifications of the Dutch Book Argument. These modifications sprung to my mind by taking a mechanism design perspective in which one is free to *create*. Firstly, a mantra in non-formal epistemology is: 'seek truth and avoid error'. There is a natural sense in which credences with $b_F = 0$ whenever $2|F| < |\Omega|$ and $b_F = 1$ whenever $2|F| > |\Omega|$ best comply with the mantra. In Appendix A, these credences as well as the credences complying with the Principle of Indifference are shown to be following from modifications of de Finetti's argument.

An assessment (of the epistemic force) of all these Dutch Book Arguments is outside the scope of this paper. What we think about these arguments turns, in part, on the design of the games. While an assessment is outside the scope, it is, however, clear that different writers will have (very) different assessments.

De Finetti also wanted his argument to elicit an agent's degrees of belief. In his later writings, de Finetti turned against the game-theoretic approach noting that a bettor agent playing his game will have some ideas as to how the book-maker will chose stakes. These

³ What exactly constitutes rational behaviour in games has become a field of study in its own right, see, e.g., Binmore (1987).

ideas inform the bettor's betting rates which raises the pragmatic worry that operationalising the Dutch Book Argument to elicit probabilistic credences from the bettor may fail, see de Finetti (1974, Section 3.6.3). The designer of the game, de Finetti, did hence not satisfy his own preferences by designing the Dutch Book Argument.

3.3 Epistemic Utility Theory

Unsatisfied with his Dutch Book Argument, de Finetti (1974) gave a new argument for Probabilism. It only concerns one agent and his hence not susceptible to strategic game-theoretic considerations. This argument was made technically more sophisticated in Joyce (1998) and Predd et al. (2009). The innovative approach walked by Joyce is to construe utilities as 'epistemic' utilities. The resulting argument is hence an epistemic, i.e., non-pragmatic, argument for Probabilism. This led to the creation of a new region of formal epistemology, epistemic utility theory, see further Pettigrew (2013b; 2016).

Pettigrew (2013b; 2013a; 2015; 2016) charted Joyce's single-agent argument for Probabilism as follows:

- (A) *Rules* On a finite set of possible worlds, Ω , containing the actual world $\omega^* \in \Omega$, fix a *scoring rule*⁴ within a particular class of scoring rules. The scoring rule measures epistemic utility of credences. The agent has to assign a credence to all $F \subseteq \Omega$, $c_F \in \mathbb{R}$.
- (B) *Outcome* At the world $\omega \in \Omega$ the score of the assigned credences.
- (C) *Utility* The agent's utility at world ω is the score of the assigned credences c .
- (D) *Decision norm* The agent ought to avoid dominated credences, do not adopt any credences c such that there exist some other credences c' which have a better score at all possible worlds. All credences which are not score-dominated are admissible, it is admissible to adopt c , if for all other credences c' there exists some world ω at which c has a strictly better score than c' .
- (E) *Mathematical theorem* All non-probabilistic credences are score-dominated. Furthermore, no probabilistic credences are score-dominated.

Following this map, Pettigrew and others have devised epistemic justifications of a number of well-known principles of rational belief formation: Bayesianism (Leitgeb and Pettigrew 2010a; 2010b), see Pettigrew (2016b) for the Principle of Indifference, Pettigrew (2012; 2013a) for the Principal Principle, Probabilism (Joyce 2009), Probabilism for a many-valued logic (Janda 2016), and imprecise probabilities (Konek 2016).

It is a well-kept secret, that all these philosophers have been engaged in one-round single-player mechanism design. That they have indeed engaged in this enterprise, should be apparent to the reader. It is perhaps not surprising that the connection to game theory has not been flagged up before, after all, epistemic utility theory was based on a move *away* from game theoretic territory. So, while epistemic utility theory provides a new interpretation it continues to make use of the argument structure of de Finetti's Dutch Book delineated in Sect. 2.3.

⁴ Scoring rules feature heavily in statistics for assessing forecasters and elicitation of probabilistic beliefs, see Gneiting and Raftery (2007) for an overview. In the current setting they are employed for belief *formation*. Note that it is a consequence of the approach that formed beliefs are probabilistic and not a stipulation as it is in statistical texts, for further discussion on these differences see Landes (2015).

An assessment (of the epistemic force) of all these epistemic utility arguments is outside the scope of this paper. What we think about these arguments turns, in part, on the design of the games. In particular, there are disagreements about the choice of a/the appropriate scoring rule, e.g., Pettigrew (2016), McCutcheon (2019). While an assessment of which scoring rule or rules is/are appropriate is outside the scope of this paper, it is, however, clear that different writers will have (very) different assessments.

3.4 Objective Bayesian Epistemology

Objective Bayesians—as opposed to subjective Bayesians—are not satisfied by adopting any old probabilistic credence function. They claim that the holy grail credences are given by the probability function which is calibrated to the evidence and has maximum entropy otherwise (Williamson 2010; 2022). In other words, objective Bayesians apply the Maximum Entropy Principle. See Paris and Vencovská (1989; 1990; 1997) for the classical justifications for the maximum entropy principle in terms of so-called ‘common sense principles’. Other classic works are Shore and Johnson (1980), Jaynes (2003), Shannon (1948), Rosenkrantz (1977). Recent works on objective Bayesianism consider multi-agent settings (Wilmers 2015), infinite domains (Landes et al. 2022; 2021a; Landes 2021a; 2021c; Williamson 2017) and computational aspects (Landes and Williamson 2016; 2022).

Recent works building on Topsøe (1979); see for example Grünwald and Dawid (2004), Landes and Williamson (2013), Landes and Williamson (2015), also argue for this seemingly outlandish claim. This work can be mapped out as a one-round two-player game between an agent and nature:

- (A) *Rules* On a finite set of possible worlds, Ω , containing the actual world, $\omega^* \in \Omega$, fix a set of probability functions defined on Ω , \mathbb{E} , and also fix a logarithmic scoring rule. The agent has to assign a credence to all $F \subseteq \Omega$, $d_F \in \mathbb{R}$ – or to all worlds $\omega \in \Omega$. Afterwards, nature chooses some probability function in \mathbb{E} .
- (B) *Outcome* The credences chosen by the agent and nature’s choice of a probability function.
- (C) *Utility* The agent receives the expected logarithmic score of the assigned credences, where expectations are taken with respect to nature’s choice in \mathbb{E} .
- (D) *Decision norm* The agent ought to adopt credences with maximal worst-case expected logarithmic score.
- (E) *Mathematical theorem* The credences which are calibrated to the evidence; i.e., the credences have to be in \mathbb{E} ; and which have maximum entropy otherwise are the only credences with maximal worst-case expected logarithmic score.

While such a map for objective Bayesianism has never been drawn before, objective Bayesians openly state that they are using game theory in their argumentation.

An assessment (of the epistemic force) of all these arguments from objective Bayesians is outside the scope of this paper. What we think about these arguments turns, in part, on the design of the games. In particular, there are disagreements about the implementation of a logarithmic scoring rule (Landes and Williamson 2013; Landes and Masterton 2017; Crupi et al. 2018). While an assessment is outside the scope of this paper, it is, however, clear that different writers will have (very) different assessments.

While all the three types of arguments in formal epistemology employ rather similar games (agents adopting credences), they stipulate different utilities and appeal to different

decision norms (sure loss avoidance, avoidance of dominated credence and maximising worst-case expected utility). Although, the arguments were devised to support *different* norms of rationality, all arguments can be construed as epistemological mechanism design.

4 Mutual Relevance

One might wonder why mechanism design and formal epistemology are mutually relevant; one may think that formal epistemology is concerned with the inescapable aspects of cognitive life, rational belief formation. On the other hand, mechanisms are designed and implemented with the sole purpose of furthering the designer's economic (or social) aims. Designed mechanisms are hence easily avoided. And so, the beautifully crafted mechanisms lack relevance for formal epistemology.

Here are some reasons to think that these fields are mutually relevant.

- (i) While it is true that market actors have great freedom to choose their trading partners, it is not always the case that a market actor can avoid transactions with the designer of a mechanism. For example, those who require a monopolist's goods have no choice; they have to trade with the monopolist. The mobile phone service providers discussed in Sect. 2.4 must take part in licensing auctions designed by states.
- (ii) Clever epistemic agents may elect not to play games featuring in Dutch Book Arguments on the grounds that they promise them very little in return. Although avoidable in principle, Dutch Book Arguments have been hugely influential. So, formal epistemology is also concerned with games which may be avoidable to some degree.
- (iii) As we saw in Sect. 3.2 on Dutch Book Arguments there exists a number of different such arguments set in in different situations which have fundamentally opposed purported rational credences. This makes one question whether there really is *one single* game against nature a rational agent cannot avoid, that is so important, that credences must be optimal for this particular game and (possibly) very bad for other games. Indeed, taking a mechanism design perspective, naturally lead to the discovery of DBAs detailed in Appendix A.⁵ But if there is no such game, then formal epistemology and mechanism design are also close in spirit and hence mutually relevant.
- (iv) Epistemologists are part of a publish-or-perish environment in which they 'have to' devise novel justifications of rationality norms. The designed game (including rules, outcomes, utilities and the decision norm) becomes almost a free parameter. The philosophical value of the designed game is then assessed in terms of how well the designed game captures rationality and thus furthers the epistemologist's career. A designed mechanism is—from the outset—a free parameter. The economic (or social) value of the designed mechanism is assessed in terms of how strongly it furthers the designer's aims. The parallels between the epistemologists and the designers of mechanisms are stark, the mutual relevance of the sub-fields follows.

⁵ Taking such a perspective may be helpful for other philosophical problems, too. For example, there is no consensus on the actions of a rational agent in the prisoner's dilemma, although (because?) the philosophical literature focuses almost exclusively on the perspective of the incarcerated agent (Kuhn 2017). But why are we so obsessed with this dilemma *designed* by the minister of justice which fails to satisfy the minister's preferences (confessing prisoners)? There seems to be ample room for designing mechanisms which will better satisfy the minister's preferences.

5 Conclusions

This paper laid bare hitherto unnoticed analogies between mechanism design and formal epistemology. The study of computer science and economics literature may thus be helpful for, say, our understanding of boundedly rational agents (Simon 1959) and/or enlarging the area of applicability (of justifications) of norms of rationality. The expertise of mechanism designers can significantly advance our grasp of concepts and techniques in formal epistemology.

For example, economists are well-aware that economic agents bend and break rules at times. Such behaviour renders all the mathematical theorems moot. Designers of auctions have investigated mechanisms for dealing with such shady agents (Klemperer 2002; Laffont and Martimort 2000). Agents which can bend the rules are only just entering the philosophical literature, see Greaves (2013) for shady epistemic agents and Konek and Levinstein (2019) for a first proposal how to deal with these agents. It seems sensible to hope that such shared problems possess shared solutions.

Furthermore, philosophers are well-equipped to critically appraise current work in economics and computer science, after all they have experience designing mechanisms.

A Proofs

We design two games G_1, G_2 played between a bettor and a book-maker. Both games differ from de Finetti's game in that the book-maker can bring about the state of the world, the bettor plays against God. The decision norm is hence to adopt credences with maximum worst-case utility. G_1 also modifies the rules of de Finetti's game, the bettor is only allowed to assign probabilistic credences.

Here is the two-player, one-round, zero-sum game G_1

- *Rules* On a finite set of possible worlds, Ω , containing the actual world $\omega^* \in \Omega$, the bettor has to assign a betting rate to all $F \subseteq \Omega$, $b_F \in \mathbb{R}$, which satisfy the axioms of probability. The book-maker chooses a stake $s_F \in \{-1, 0, +1\}$ for every $F \subseteq \Omega$; selling or buying the bet at the betting rate b_F . The book-maker also chooses which possible world $\omega \in \Omega$ obtains.
- *Outcome* The probabilistic betting rates b , the stakes s and the actual world ω^* brought about by the book-maker.
- *Utility* At the world $\omega \in \Omega$, the bet on F returns $r_\omega(b_F) = -s_F \cdot b_F$, if $\omega \notin F$ and $r_\omega(b_F) = s_F \cdot (1 - b_F)$, if $\omega \in F$. The bettor's utility at world ω is the return of all bets, $\sum_{F \subseteq \Omega} r_\omega(b_F)$. The book-maker's utility is $-\sum_{F \subseteq \Omega} r_\omega(b_F)$.
- *Decision norm* The bettor ought to maximise worst-case utility, adopt credences b such that $\min_{\omega \in \Omega} \min_{s_F} \sum_{F \subseteq \Omega} r_\omega(b_F)$ is maximal.
- *Mathematical theorem* The bettor maximises worst-case utility, if and only if $b_F(\omega) = \frac{1}{|\Omega|}$.

Here is a proof of the theorem.

Proof For a fixed possible world $\omega \in \Omega$ the bookmaker will chose stake $s_F = -1$, if $\omega \notin F$ and $s_F = 1$, if $\omega \in F$, regardless of the betting rates. The agent's utility at world $\omega \in \Omega$ is hence

$$\begin{aligned}
 \sum_{F \subseteq \Omega} r_{\omega}(b_F) &= \sum_{\substack{\emptyset \subset F \subseteq \Omega \\ \omega \in F}} b_F + \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} (1 - b_F) \\
 &= \sum_{\substack{\emptyset \subset F \subseteq \Omega \\ \omega \in F}} b_F - \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} b_F + \left(\frac{2^{|\Omega|}}{2}\right).
 \end{aligned}
 \tag{1}$$

To maximise this expression we can ignore the final term, $\frac{2^{|\Omega|}}{2}$. The agent will hence obtain maximum utility, if and only if

$$\begin{aligned}
 \min_{\omega \in \Omega} \sum_{\substack{\emptyset \subset F \subseteq \Omega \\ \omega \in F}} b_F - \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} b_F &= \min_{\omega \in \Omega} \sum_{\substack{\emptyset \subset F \subseteq \Omega \\ \omega \in F}} b_{\omega} + b_{F \setminus \{\omega\}} - \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} b_F \\
 &= \min_{\omega \in \Omega} \frac{2^{|\Omega|}}{2} b_{\omega} + \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} b_F - \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} b_F \\
 &= \min_{\omega \in \Omega} \frac{2^{|\Omega|}}{2} b_{\omega}
 \end{aligned}$$

is maximal. It is maximal, if and only if $b_{\omega} = \frac{1}{|\Omega|}$ for all $\omega \in \Omega$ (recall that we only consider probabilistic credences here). □

Here is game G_2 :

- *Rules* On a finite set of possible worlds, Ω , containing the actual world $\omega^* \in \Omega$, the bettor has to assign a betting rate to all $F \subseteq \Omega$, $b_F \in \mathbb{R}$. The book-maker chooses a stake $s_F \in \{-1, 0, +1\}$ for every $F \subseteq \Omega$; selling or buying the bet at the betting rate b_F . The book-maker also chooses which possible world $\omega \in \Omega$ obtains.
- *Outcome* The betting rates b , the stakes s and the actual world ω^* brought about by the book-maker.
- *Utility* At the world $\omega \in \Omega$, the bet on F returns $r_{\omega}(b_F) = -s_F \cdot b_F$, if $\omega \notin F$ and $r_{\omega}(b_F) = s_F \cdot (1 - b_F)$, if $\omega \in F$. The bettor's utility at world ω is the return of all bets, $\sum_{F \subseteq \Omega} r_{\omega}(b_F)$. The book-maker's utility is $-\sum_{F \subseteq \Omega} r_{\omega}(b_F)$.
- *Decision norm* The bettor ought to maximise worst-case utility, adopt credences b such that $\min_{\omega \in \Omega} \min_{s_F} \sum_{F \subseteq \Omega} r_{\omega}(b_F)$ is maximal.
- *Mathematical theorem* The bettor maximises worst-case utility, if and only if there exists some constant $c \in [0, 1]$ such that

$$b_F = \begin{cases} 1, & \text{if } 2|F| > |\Omega| \\ c, & \text{if } 2|F| = |\Omega| \\ 0, & \text{if } 2|F| < |\Omega|. \end{cases}$$

We now give a proof of the theorem.

Proof The bettor's utility is again given by (1). Suppose that b is not invariant under permutations of the elements of Ω . Let b' be the symmetrisation of b , the arithmetic mean

of all credences which can be obtained from b by permuting elements of Ω . Note that $\min_{\omega \in \Omega} \min_{s_F} \sum_{F \subseteq \Omega} r_{\omega}(b_F) < \min_{\omega \in \Omega} \min_{s_F} \sum_{F \subseteq \Omega} r_{\omega}(b'_F)$.

We shall thus focus on symmetric credences, i.e., b_F only depends on $|F|$ and write b_m for $m = |F|$. We find

$$\begin{aligned} \sum_{\substack{\emptyset \subset F \subseteq \Omega \\ \omega \in F}} b_F - \sum_{\substack{\emptyset \subseteq F \subset \Omega \\ \omega \notin F}} b_F &= \sum_{m=1}^{|\Omega|} \binom{|\Omega| - 1}{m - 1} b_m - \sum_{m=0}^{|\Omega| - 1} \binom{|\Omega| - 1}{m} b_m \\ &= b_{\Omega} - b_{\emptyset} + \sum_{m=1}^{|\Omega| - 1} \left(\binom{|\Omega| - 1}{m} - \binom{|\Omega| - 1}{m - 1} \right) b_m \\ &= b_{\Omega} - b_{\emptyset} + \sum_{m=1}^{|\Omega| - 1} \frac{(|\Omega| - 1)!}{(m - 1)! (|\Omega| - m - 1)!} \left(\frac{1}{m} - \frac{1}{|\Omega| - m} \right) b_m \\ &= b_{\Omega} - b_{\emptyset} + \sum_{m=1}^{|\Omega| - 1} \frac{(|\Omega| - 1)!}{m! (|\Omega| - m)!} (|\Omega| - 2m) b_m. \end{aligned}$$

Clearly, this is maximal, if and only if

$$b_F = \begin{cases} 1, & \text{if } 2|F| > |\Omega| \\ 0, & \text{if } 2|F| < |\Omega|. \end{cases}$$

Hence, the bettor maximises worst-case utility, if and only if there exists some constant $c \in [0, 1]$ such that

$$b_F = \begin{cases} 1, & \text{if } 2|F| > |\Omega| \\ c, & \text{if } 2|F| = |\Omega| \\ 0, & \text{if } 2|F| < |\Omega|. \end{cases}$$

□

I’m indebted to Jeff Paris who communicated these proofs privately to me, after I had conjectured something half-baked.

Acknowledgements Many thanks to Teddy Groves, Hannes Leitgeb, Barbara Osimani, Jeff Paris, Mantas Radzvilas, Roland Poellinger, Reuben Stern and Jon Williamson for providing helpful comments. I gratefully acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) 405961989 and 432308570, the European Research Council (Grant 639276) and NextGenerationEU funding for the project “Practical Reasoning for Human-Centred Artificial Intelligence”.

Funding Open access funding provided by Università degli Studi di Milano within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abrache, J., T. G. Crainic, M. Gendreau, and M. Rezik. 2007. Combinatorial auctions. *Annals of OR* 153 (1): 131–164. <https://doi.org/10.1007/s10479-007-0179-z>.
- Binmore, K. 1987. Modeling rational players: Part I. *Economics and Philosophy* 3 (2): 179–214. <https://doi.org/10.1017/S0266267100002893>.
- Bradley, S. 2015. Imprecise Probabilities. In *Stanford encyclopedia of philosophy*, ed. E. N. Zalta, Summer 2015 edition. <http://plato.stanford.edu/entries/imprecise-probabilities>.
- Branco, F. 1997. The design of multidimensional auctions. *The Rand journal of economics* 28 (1): 63–81. <http://www.jstor.org/stable/2555940>.
- Buettner, R., and J. Landes. 2012. Web service based applications for electronic labor markets: A multi-dimensional price VCG auction with individual utilities. In *Proceedings of ICIW*, 168–177. http://www.thinkmind.org/index.php?view=article&articleid=icw_2012_6_20_20045.
- Chawla, S., and B. Sivan. 2014. Bayesian algorithmic mechanism design. *SIGecom Exchange* 13 (1): 5–49. <https://doi.org/10.1145/2692375.2692378>.
- Che, Y. 1993. Design competition through multidimensional auctions. *Rand Journal of Economics* 24 (4): 668–680. <http://www.jstor.org/stable/2555752>.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11 (1): 17–33. <https://doi.org/10.1007/BF01726210>.
- Cramton, P. 2003. Electricity market design: the good, the bad, and the ugly. In *Proceedings of HICSS*, 8. <https://doi.org/10.1109/HICSS.2003.1173866>.
- Cramton, P., Y. Shoham, and R. Steinberg. 2006. *Combinatorial Auctions*. Cambridge: MIT Press.
- Crupi, V., J. Nelson, B. Meder, et al. 2018. Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science* 42: 1410–1456. <https://doi.org/10.1111/cogs.12613>.
- Dash, R. K., N. R. Jennings, and D. C. Parkes. 2003. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems* 18: 40–47. <https://doi.org/10.1109/MIS.2003.1249168>.
- de Finetti, B. 1937. La prévision: ses lois logiques, ses sources subjectives. In *Annales de l'institut Henri Poincaré*, volume 7, 1–68. [http://archive.numdam.org/ARCHIVE/AIHP/AIHP_1937__7_1/AIHP_1937__7_1_1_0.pdf](http://archive.numdam.org/ARCHIVE/AIHP/AIHP_1937__7_1/AIHP_1937__7_1_1_0/AIHP_1937__7_1_1_0.pdf).
- de Finetti, B. 1974. *Theory of Probability*. Chichester: Wiley.
- de Finetti, B. 1980. Foresight: Its logical laws, its subjective sources. In *Studies in Subjective Probability*, eds. H. E. Kyburg and H. E. Smokler, 53–118. Huntington: Krieger, 2 edition.
- Decock, L., I. Douven, and M. Sznajder. 2016. A geometric principle of indifference. *Journal of Applied Logic* 19: 54–70. <https://doi.org/10.1016/j.jal.2016.05.002>.
- Easwaran, K. 2011a. Bayesianism I: introduction and arguments in favor. *Philosophy Compass* 6 (5): 312–320. <https://doi.org/10.1111/j.1747-9991.2011.00399.x>.
- Easwaran, K. 2011b. Bayesianism II: applications and criticisms. *Philosophy Compass* 6 (5): 321–332. <https://doi.org/10.1111/j.1747-9991.2011.00398.x>.
- Eva, B. 2019. Principles of indifference. *The Journal of Philosophy* 116 (7): 390–411. <https://doi.org/10.5840/jphil2019116724>.
- Gneiting, T., and A. E. Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102 (477): 359–378. <https://doi.org/10.1198/016214506000001437>.
- Greaves, H. 2013. Epistemic decision theory. *Mind* 122 (488): 915–952. <https://doi.org/10.1093/mind/fzt090>.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41 (4): 617–631. <http://bbs.cenet.org.cn/UploadImages/200642013305286207.pdf>.
- Grünwald, P. D., and A. P. Dawid. 2004. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Annals of Statistics* 32 (4): 1367–1433. <https://doi.org/10.1214/00905360400000553>.
- Gyenis, B., and L. Wroński. 2017. *Is it the Principal Principle that Implies the Principle of Indifference?*, 35–41. Cham: Springer. https://doi.org/10.1007/978-3-319-55486-0_2.
- Hájek, A. 2008. Arguments for - or against - Probabilism? *British Journal for the Philosophy of Science* 59 (4): 793–819. <https://doi.org/10.1093/bjps/axn045>.
- Hartline, J. D. 2012. Bayesian mechanism design. *Foundations and Trends in Theoretical Computer Science* 8 (3): 143–263. <https://doi.org/10.1561/04000000045>.
- Hawthorne, J., J. Landes, C. Wallmann, and J. Williamson. 2017. The principal principle implies the principle of indifference. *British Journal for the Philosophy of Science* 68: 123–131. <https://doi.org/10.1093/bjps/axv030>.

- Hohner, G., J. Rich, E. Ng, et al. 2003. Combinatorial and quantity-discount procurement auctions benefit mars, incorporated and its suppliers. *Interfaces* 33 (1): 23–35. <https://doi.org/10.1287/inte.33.1.23.12717>.
- Houser, D., and J. Wooders. 2006. Reputation in auctions: theory, and evidence from eBay. *Journal of Economics & Management Strategy* 15 (2): 353–369. <https://doi.org/10.1111/j.1530-9134.2006.00103.x>.
- Huber, F. 2016. Formal representations of belief. In *Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, spring 2016 edition. <https://plato.stanford.edu/archives/spr2016/entries/formal-belief/>.
- Janda, P. 2016. Measuring inaccuracy of uncertain doxastic states in many-valued logical systems. *Journal of Applied Logic* 14: 95–112. <https://doi.org/10.1016/j.jal.2015.09.010>.
- Jaynes, E. T. 2003. *Probability Theory: The Logic of Science*. Cambridge: Cambridge University Press.
- Jehiel, P., B. Moldovanu, and E. Stacchetti. 1996. How (not) to sell nuclear weapons. *American Economic Review* 86 (4): 814–829. <http://www.jstor.org/stable/2118306>.
- Joyce, J. M. 1998. A nonpragmatic vindication of probabilism. *Philosophy of Science* 65 (4): 575–603. <https://doi.org/10.1086/392661>.
- Joyce, J. M. 2009. Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In *Degrees of Belief*, eds. F. Huber and C. Schmidt-Petri, volume 342 of *Synthese Library*, 263–297. Dordrecht: Springer. https://doi.org/10.1007/978-1-4020-9198-8_11.
- Joyce, J. M. 2011. The Development of Subjective Bayesianism. In *Handbook of the History of Logic. Vol. 10: Inductive Logic*, eds. D. Gabbay, S. Hartmann, and J. Woods, 415–475. Oxford: North Holland.
- Konek, J. 2016. Epistemic conservativity and imprecise credence. *forthcoming*. <https://dl.dropboxusercontent.com/s/6yl8sx0rkk0cey/ecic.pdf>.
- Klemperer, P. 1999. Auction theory: A guide to the literature. *Journal of Economic Surveys* 13 (3): 227–286. <https://doi.org/10.1111/1467-6419.00083>.
- Klemperer, P. 2002. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review* 46 (4–5): 829–845. [https://doi.org/10.1016/S0014-2921\(01\)00218-5](https://doi.org/10.1016/S0014-2921(01)00218-5).
- Konek, J., and B. Levinstein. 2019. The foundations of epistemic decision theory. *Mind* 128: 69–107. <https://doi.org/10.1093/mind/fzw044>.
- Kopec, M. 2017. A new group Dutch book argument. *Ratio* 30 (2): 122–136. <https://doi.org/10.1111/rati.12131>.
- Kuhn, S. 2017. Prisoner’s Dilemma. In *Stanford encyclopedia of philosophy*, ed. E. N. Zalta, Winter 2019 edition. <http://plato.stanford.edu/archives/fall2014/entries/prisoner-dilemma/>.
- Laffont, J.-J., and D. Martimort. 2000. Mechanism design with collusion and correlation. *Econometrica* 68 (2): 309–342. <https://doi.org/10.1111/1468-0262.00111>.
- Landes, J. 2015. Probabilism, entropies and strictly proper scoring rules. *International Journal of Approximate Reasoning* 63: 1–21. <https://doi.org/10.1016/j.ijar.2015.05.007>.
- Landes, J. 2021a. A triple uniqueness of the maximum entropy approach. In *Proceedings of ECSQARU*, eds. J. Vejnarová and N. Wilson, volume 12897 of *LNAI*, 644–656, Cham: Springer. https://doi.org/10.1007/978-3-030-86772-0_46.
- Landes, J. 2021b. Bayesian epistemology. In *Oxford bibliographies in philosophy*. New York: Oxford University Press. <https://doi.org/10.1093/obo/9780195396577-0417>.
- Landes, J. 2021c. The Entropy-limit (Conjecture) for Σ_2 -Premisses. *Studia Logica* 109: 423–442. <https://doi.org/10.1007/s11225-020-09912-3>.
- Landes, J., and G. Masterton. 2017. Invariant equivocation. *Erkenntnis* 82: 141–167. <https://doi.org/10.1007/s10670-016-9810-1>.
- Landes, J., and J. Williamson. 2013. Objective Bayesianism and the maximum entropy principle. *Entropy* 15 (9): 3528–3591. <https://doi.org/10.3390/e15093528>.
- Landes, J., and J. Williamson. 2015. Justifying objective Bayesianism on predicate languages. *Entropy* 17 (4): 2459–2543. <https://doi.org/10.3390/e17042459>.
- Landes, J., and J. Williamson. 2016. Objective Bayesian nets from consistent datasets. In *Proceedings of MaxEnt*, eds. A. Giffin and K. H. Knuth, volume 1757. <https://doi.org/10.1063/1.4959048>.
- Landes, J., and J. Williamson. 2022. Objective Bayesian nets for integrating consistent datasets. *Journal of Artificial Intelligence Research* 74: 393–458. <https://doi.org/10.1613/jair.1.13363>.
- Landes, J., S. Rafiee Rad, and J. Williamson. 2021a. Towards the entropy-limit conjecture. *Annals of Pure and Applied Logic*, 172. <https://doi.org/10.1016/j.apal.2020.102870>.
- Landes, J., C. Wallmann, and J. Williamson. 2021b. The Principal Principle, admissibility, and normal informal standards of what is reasonable. *European Journal for Philosophy of Science*, 11. <https://doi.org/10.1007/s13194-021-00352-x>.
- Landes, J., S. Rafiee Rad, and J. Williamson. 2022. Determining maximal entropy functions for objective Bayesian inductive logic. *Journal of Philosophical Logic*. <https://doi.org/10.1007/s10992-022-09680-6>.

- Leitgeb, H., and R. Pettigrew. 2010a. An objective justification of Bayesianism I: measuring inaccuracy. *Philosophy of Science* 77 (2): 201–235. <https://doi.org/10.1086/651317>.
- Leitgeb, H., and R. Pettigrew. 2010b. An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science* 77 (2): 236–272. <https://doi.org/10.1086/651318>.
- Lewis, D. 1980. A Subjectivist's Guide to Objective Chance. In *Studies in inductive logic and probability*, ed. R. C. Jeffrey, volume 2, 263–293. Berkeley: University of California Press.
- Malone, T. W., J. Yates, and R. I. Benjamin. 1987. Electronic markets and electronic hierarchies. *Communications of the ACM* 30 (6): 484–497. <https://doi.org/10.1145/214762.214766>.
- McCutcheon, R. G. 2019. In favor of logarithmic scoring. *Philosophy of Science* 86 (2): 286–303. <https://doi.org/10.1086/702028>.
- Nisan, N. 2014. Algorithmic mechanism design. In *Handbook of game theory*, 478–514. Amsterdam: Elsevier.
- Novack, G. 2010. A defense of the principle of indifference. *Journal of Philosophical Logic* 39 (6): 655–678. <https://doi.org/10.1007/s10992-010-9147-1>.
- Ortner, R., and H. Leitgeb. 2011. Mechanizing induction. In *Handbook of the history of logic*, eds. D. M. Gabbay, S. Hartmann and J. Woods, 719–772. <https://doi.org/10.1016/b978-0-444-52936-7.50018-5>.
- Paris, J. B. 2001. A note on the Dutch Book method. In *Proceedings of ISIPTA*, eds. G. De Cooman, T. Fine and T. Seidenfeld, 301–306.
- Paris, J. B. 2014. What you see is what you get. *Entropy* 16 (11): 6186–6194. <https://doi.org/10.3390/e16116186>.
- Paris, J. B., and A. Vencovská. 1989. On the applicability of maximum entropy to inexact reasoning. *International Journal of Approximate Reasoning* 3 (1): 1–34. [https://doi.org/10.1016/0888-613X\(89\)90012-1](https://doi.org/10.1016/0888-613X(89)90012-1).
- Paris, J. B., and A. Vencovská. 1990. A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning* 4 (3): 183–223. [https://doi.org/10.1016/0888-613X\(90\)90020-3](https://doi.org/10.1016/0888-613X(90)90020-3).
- Paris, J. B., and A. Vencovská. 1997. In defense of the maximum entropy inference process. *International Journal of Approximate Reasoning* 17 (1): 77–103. [https://doi.org/10.1016/S0888-613X\(97\)00014-5](https://doi.org/10.1016/S0888-613X(97)00014-5).
- Pettigrew, R. 2012. Accuracy, chance, and the principal principle. *The Philosophical Review* 121 (2): 241–275. <https://doi.org/10.1215/00318108-1539098>.
- Pettigrew, R. 2013a. A new epistemic utility argument for the principal principle. *Episteme* 10 (1): 19–35. <https://doi.org/10.1017/epi.2013.5>.
- Pettigrew, R. 2013b. Epistemic utility and norms for Credences. *Philosophy Compass* 8 (10): 897–908. <https://doi.org/10.1111/phc3.12079>.
- Pettigrew, R. 2015. Interview with Richard Pettigrew. *Reasoner* 9 (4): 28–31 (<http://blogs.kent.ac.uk/thereasoner/files/2015/04/TheReasoner-94.pdf>).
- Pettigrew, R. 2016a. *Accuracy and the laws of credence*. Oxford: Oxford University Press.
- Pettigrew, R. 2016b. Accuracy, risk, and the principle of indifference. *Philosophy and Phenomenological Research*, 92 (1): 35–59. <https://doi.org/10.1111/phpr.12097>.
- Pettigrew, R. 2020a. The principal principle does not imply the principle of indifference. *British Journal for the Philosophy of Science* 71: 605–619. <https://doi.org/10.1093/bjps/axx060>.
- Pettigrew, R. 2020b. *Dutch book arguments*. Elements in Decision Theory and Philosophy. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108581813>.
- Predd, J., R. Seiringer, E. Lieb, et al. 2009. Probabilistic coherence and proper scoring rules. *IEEE Transactions on Information Theory* 55 (10): 4786–4792. <https://doi.org/10.1109/TIT.2009.2027573>.
- Radzvilas, M., W. Peden, and F. D. Pretis. 2022. A Battle in the Statistics Wars: a simulation-based comparison of Bayesian. *Frequentist and williamsonian methodologies*. *Synthese*. <https://doi.org/10.1007/s11229-021-03395-y>.
- Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood. 2006. The value of reputation on eBay: A controlled experiment. *Experimental Economics* 9: 79–101. <https://doi.org/10.1007/s10683-006-4309-2>.
- Rosenkrantz, R. D. 1977. *Inference, method and decision: towards a Bayesian philosophy of science*. Dordrecht: Reidel.
- Roth, A. E. 2002. The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica* 70 (4): 1341–1378. <https://doi.org/10.1111/1468-0262.00335>.
- Sandholm, T. 2007. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine* 28 (3): 45–58. <https://doi.org/10.1609/aimag.v28i3.2054>.
- Schupbach, J. N. 2022. *Bayesianism and scientific reasoning*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108657563>.
- Shafer, G. 2011. A betting interpretation for probabilities and Dempster-Shafer degrees of belief. *International Journal of Approximate Reasoning* 52 (2): 127–136. <https://doi.org/10.1016/j.ijar.2009.05.012>.

- Shafer, G., and V. Vovk. 2001. *Probability and finance*. New York: Wiley.
- Shannon, C. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27: 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.
- Shore, J. E., and R. W. Johnson. 1980. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory* 26 (1): 26–37. <https://doi.org/10.1109/TIT.1980.1056144>.
- Simon, H. A. 1959. Theories of decision-making in economics and behavioral science. *The American Economic Review* 49 (3): 253–283. <http://www.jstor.org/stable/1809901>.
- Skipper, M., and J. C. Bjerring. 2022. Bayesianism for non-ideal agents. *Erkenntnis* 87 (1): 93–115. <https://doi.org/10.1007/s10670-019-00186-3>.
- Spohn, W. 2012. *The laws of belief*. Ranking Theory and its Philosophical Applications. Oxford: Oxford University Press.
- Sprenger, J., and S. Hartmann. 2019. *Bayesian philosophy of science*. Oxford: Oxford University Press.
- Titelbaum, M. G., and C. Hart. 2020. The principal principle does not imply the principle of indifference, because conditioning on biconditionals is counterintuitive. *British Journal for the Philosophy of Science* 71. <https://doi.org/10.1093/bjps/axy011>.
- Topsøe, F. 1979. Information theoretical optimization techniques. *Kybernetika* 15: 1–27 (<http://eudml.org/doc/27475>).
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (1): 8–37. <http://www.jstor.org/stable/2977633>.
- Weatherston, B. 2003. From classical to intuitionistic probability. *Notre Dame Journal of Formal Logic* 44 (2): 111–123. <https://doi.org/10.1305/ndjfl/1082637807>.
- Weisberg, J. 2015. You've come a long way, Bayesians. *Journal of Philosophical Logic* 44 (6): 817–834. <https://doi.org/10.1007/s10992-015-9363-9>.
- White, R. 2010. Evidential symmetry and mushy credence. In *Oxford studies in epistemology*, ed. T.S. Gendler and J. Hawthorne, volume 3, 161–186. Oxford: Oxford University Press.
- Williamson, J. 2004. A dynamic interaction between machine learning and the philosophy of science. *Minds and Machines* 14 (4): 539–549. <https://doi.org/10.1023/b:mind.0000045990.57744.2b>.
- Williamson, J. 2010. *In defence of objective Bayesianism*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199228003.001.0001>.
- Williamson, J. 2017. *Lectures on inductive logic*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199666478.001.0001>.
- Williamson, J. 2022. Direct inference and probabilistic accounts of induction. *Journal for General Philosophy of Science*. <https://doi.org/10.1007/s10838-021-09584-0>.
- Wilmers, G. 2015. A foundational approach to generalising the maximum entropy inference process to the multi-agent context. *Entropy* 17 (2): 594–645. <https://doi.org/10.3390/e17020594>.

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