

(The final version is due to be published in *Mind and Language*. Please refer to the published version.)

Should credence be sensitive to practical factors? A cost-benefit analysis

Jie Gao

Zhejiang University

According to evidentialist views, credence in a proposition p should be proportional to the degree of evidential support that one has in favor of p . However, empirical evidence suggests that our credences are systematically sensitive to practical factors. In this paper, I provide a cost-benefit analysis of credences' practical sensitivity. The upshot of this analysis is that credences sensitive to practical factors fare better than practically insensitive ones along several dimensions. All things considered, our credences should be sensitive to practical factors.

1. INTRODUCTION

When we believe something, we believe it with a certain degree of confidence. Epistemologists usually characterise one's strength of confidence in terms of degrees of credence. In a standard Bayesian framework, these degrees are represented on a scale from 0 to 1, where $cr(p) = 1$ and $cr(p) = 0$ express, respectively, maximal degree of credence (absolute certainty) that p and that not- p . According to the standard evidentialist view in epistemology, credence in a proposition p should be proportional to the degree of evidential support that one has in favor of p . In this picture, degrees of credence should be probabilistically coherent and should be exclusively updated on new evidence by following conditionalization rules.¹

However, this is not how ordinary people form and maintain their doxastic attitudes. A wide body of empirical studies has shown that our confidence is systematically sensitive to non-evidential factors. Some of these studies suggest that certain factors such as high costs of being wrong have the effect of making one more cautious and less confident in one's judgment, whereas others factors such as benefits of having a settled opinion and difficulties of acquiring further evidence have opposite effects (Maysless & Kruglanski, 1987; Kruglanski & Webster, 1991; Kruglanski, Webster, & Klem, 1993; Webster, 1993). On the basis of analyses of these empirical data, Gao (2019) has defended a view according to which our credences are systematically sensitive to different types of practical factors. Other authors have defended similar theses. Clarke (2013), Greco (2015), and Norby (2015) propose that credences are distributed

¹ There are a variety of different conditionalization rules. The most prominent ones are Bayes and Jeffrey conditionalization.

over a space of possibilities that reflects the subject's interests at a given time. In this picture, practical factors can generate changes of credence by modifying the relevant set of doxastic possibilities. Nagel (2008) also suggests that we are prone to project a lower-level subjective confidence to a high-stakes subject than to her low-stakes counterpart.²

The sensitivity of credence to practical factors (henceforth, CPS) implies the violation of widely accepted evidentialist norms. It is thus unsurprising that some philosophers have deemed it epistemically irrational (e.g., Rubin 2015; Schroeder 2018). Against such standard view, Gao (2019) suggests that even though CPS falls short of ideal standards of epistemic rationality, it instantiates *bounded rationality*, a sort of non-ideal epistemic rationality.³ This paper aims to address a more general and overarching question about the normative status of CPS, namely, whether this sort of credal sensitivity, as it is naturally instantiated by human beings in normal circumstances, is a better strategy for ordinary people's credence formation, retention and update than credal insensitivity to practical factors (henceforth, CPI). The issues I am concerned with here are whether our credences should be sensitive to practical factors *all things considered*, and whether we *all things considered* should adopt cognitive strategies resulting in CPS.⁴

A methodology that is suitable for investigating the above question is *cost-benefit analysis*. As a well-developed form of applied consequentialism, cost-benefit analysis permits us to weigh the desirability of various options available to us by comparing advantages and disadvantages of each available option along several dimensions. The upshot of my analysis is that in most cases the benefits of a moderate version of CPS remarkably outweigh the benefits of CPI, while the costs of such a CPS are very limited or negligible. Thus, all things considered, our credences should be sensitive to practical factors.⁵

The significance of the current project is multifaceted. First, in the pragmatic encroachment debate, pragmatic approaches to rational credence are often quickly dismissed given the possible violation of ideally rational constraints. As this paper aims to show, such a quick dismissal is unjustified in the light of the practical value of CPS. Moreover, assuming a plausible correlation between full belief and high

² See also Moss (2018a, 2018b) for a defence of the idea that moral features can affect whether credence constitute knowledge.

³ Other authors also mention or discuss the view that rational credences depend partly on practical factors, including Stanley (2005), Armendt (2010), Clarke (2013), Norby (2015), Greco (2015), Kim (2017), Hajek & Lin (2017), Moss (2018), Jackson (2019).

⁴ In addition to CPS and CPI, in principle we could conceive hybrid CPS-CPI approaches allowing practical sensitivity of credence in some circumstances but not in others. Such hybrid approaches may be considered mere variants of CPS. Since many of the costs I will discuss below against adopting CPI also apply to hybrid approaches, I will not explicitly consider hybrid approaches in this paper.

⁵ Not all forms of credal sensitivity will have high benefits and negligible costs. We can easily imagine a radical form of credal sensitivity that would assign only extreme credences (0 and 1) in rush-oriented cases. Arguably, this would serve one's needs very badly. As I will clarify in the next section, my focus is primarily on a moderate version of CPS that approximates human natural tendencies of doxastic regulation.

credence, one may derive similar conclusions about the rationality of belief's sensitivity to practical factors, thus fostering the pragmatic encroachment program.^{6,7}

Second, the current project could shed new light on the ongoing debate over the nature of rationality—the so-called “rationality wars” (Stein, 1996; Rysiew, 2008). This debate is about how we should interpret the fact that in our reasoning we systematically deviate from ideal norms (such as rules of logic, probability and decision theory). While the Meliorists tend to interpret such deviations as signs of irrationality (e.g., Kahneman & Tversky, 1996; Stanovich, 1999), the Panglossians contend that such deviations can be normative, in the sense of providing prescriptions about what means people should use to achieve goals given their cognitive and environmental limitations (e.g., Anderson, 1990; Gigerenzer, 2008). The same debate transposes to the evaluation of CPS. This paper can be viewed as extending the latter project to strategies concerning credence formation, retention and update.

The significance of this contribution also extends to methodological issues. Empirical studies have provided good evidence of credal sensitivity to a wide range of non-evidential factors.⁸ The cost-benefit framework employed in this paper is applicable to analyses of other aspects of credal sensitivity. More generally, the current project may contribute to reviving the cost-benefit analysis approach in traditional and applied epistemology, promoting its application to a wide variety of epistemological issues, fostering a promising new line of research in this domain.

The paper is structured as follows. Section 2 introduces credal pragmatism as a descriptive theory and discusses its prospects as a normative theory. Section 3 introduces a framework suitable for the cost-benefit analysis of CPS and CPI. This framework is based on Bishop and Trout's cost-benefit analysis of reasoning strategies. Section 4 analyses the costs and benefits of, respectively, CPS and CPI along the parameters specified in Section 3 and provides a comparison. Section 5 concludes.

2. CREDAL PRAGMATISM: DESCRIPTIVE AND NORMATIVE

Credal pragmatism has been originally conceived as a descriptive thesis, concerning how *ordinary people actually form, maintain and revise their doxastic attitudes*. We can distinguish practical factors that can affect

⁶ Pragmatic encroachment is here conceived broadly, referring not only to normative epistemic conditions or concepts such as knowledge and epistemic justification, but also to non-normative cognitive attitudes and states. See Kim (2017, §6) for a taxonomy of varieties of pragmatic encroachment.

⁷ Among proponents of practical sensitivity of both credence and belief, Clarke (2013) and Greco (2015) defend credence-one views of belief according to which belief requires maximal credence, whereas Gao (2016, 2019, forthcoming) defends a version of the threshold view.

⁸ In addition to factors mentioned above, such factors include the ease (or difficulty) of recall or imagination (Tversky & Kahneman, 1974; Lichtenstein, 1978; Johnson et al., 1993), the way in which an event is described (Johnson et al., 1993; Tversky & Kahneman, 1983; Tversky & Koehler, 1994; Rottenstreich & Tversky, 1997), and emotional reactions (Slovic et al., 1982; Loewenstein et al., 2001; Rottenstreich & Hsee, 2001).

credences into two groups. On the one hand, there are factors related to the avoidance of error. These have the effect of leading one to assign more weight to the goal of not believing p if p is false. Such factors include stakes (e.g. high costs of being wrong about p , significant benefits of achieving accuracy in judgment) and the availability of further evidence. On the other hand, there are factors related to truth-acquisition. They have the effect of leading one to assign more weight to the goal of believing p when p is true. Examples include urgency, considerable benefits of having a settled opinion, remarkable costs of not having a settled view, and difficulty of acquiring further evidence (e.g. noise in the environment, unattractiveness of the task). Call the former group *caution-oriented factors* and the latter group *rush-oriented factors*.⁹ Descriptive credal pragmatism claims that, other things being equal, rush-oriented factors tend to raise one's degree of credence in propositions one takes to be more likely and lower that in propositions one takes to be less likely. More precisely, such factors have a polarizing effect leading to more extreme credences: Depending on whether a credence is above or below 0.5, these factors move the credence closer to the extremes 1 or 0. Conversely, caution-oriented factors tend to moderate one's degree of credence—leading to credences closer to the scale midpoint 0.5.

While I am open to different ways of implementing this view, in this paper I conceive CPS as a rough approximation to the natural ways in which ordinary people's confidence happens to be sensitive to non-evidential factors, as suggested by empirical studies. This version of CPS is supposed to reflect built-in psychological or natural tendencies that normal human beings already possess. Although this sensitivity might allow individual differences, it should exclude extreme patterns such as mechanisms radically polarizing credences to either 0 or 1 under rushed-oriented factors or excessively moderating credences under caution-oriented factors. Empirical studies and ordinary patterns of behaviour rule out such extreme forms of credal sensitivity as descriptively implausible.

For present purposes we can remain somewhat vague on how exactly CPS works. However, it could be helpful to illustrate the basic idea with an example:

Peanut Sandwich

An old friend will come to visit you today. You are buying some sandwiches for her in a nearby supermarket. You tried some free samples of salmon sandwiches and they don't taste like they have peanuts inside.

(Situation A) You know that your friend does not like peanuts. Basing your judgment on your taste you form a quite high confidence in the proposition that there are no peanuts in the salmon sandwich (call it p), say $Cr(p) = 0.8$.

⁹ It is worth observing that not all credal sensitivists agree on the list of practical factors affecting credence. For instance, according to Clarke (2013) and Greco (2015), only practical factors consciously recognised by the subject (e.g., salience, perceived stakes) can result in changes in the subject's credence functions.

(Situation B) You know that your friend has a serious allergy to peanuts. Basing your judgment on your taste, this time you are not very confident in p . Your credence in p is lower than in situation A, say $Cr(p) = 0.7$.

(Situation C) You know that your friend has no allergy to peanuts. Moreover, you realise that she will be at your place very soon, so you'd better quickly make up your mind, buy something and run home. In this case, you form a high confidence in p . Your credence is higher than in situation A, say $Cr(p) = 0.9$.

In Peanut Sandwich, situation A does not involve any practical factor that could modify your confidence. As a result your confidence doesn't manifest practical sensitivity. Situation B involves sensitivity of your confidence to caution-oriented factors: Stakes are very high on being wrong about whether the sandwich contains peanuts. This factor makes you less confident. Situation C involves credal sensitivity to rush-oriented factors: The urgency to make up your mind make you slightly overconfident and form a settled judgment on which you can act quickly.

Notice that in cases manifesting credal sensitivity, the relevant practical factors are supposed to be perceived by the subject rather than actual ones; otherwise it would be impossible for practical factors to exert influence over our mental states. Moreover, in order to exert influence on our confidence, these factors do not need to be perceived at a conscious level, or even believed. In Peanut Sandwich, your credence can be affected by the relevant practical factors even though you do not consciously consider such factors and do not form a doxastic attitude about them. You may perceive and become aware of such factors through affective states such as feelings of anxiety or hurry. As shown by psychological studies, in most cases of credal sensitivity the impact of perceived practical factors is immediate, non-inferential and automatic. The cognitive response to such factors is executed at a sub-personal level, caused by emotional responses or other automatic mechanisms.¹⁰

Normative credal pragmatism (henceforth NCP) further claims that, in some sense of "rational", it is rational for credences to be sensitive to certain practical factors. Other things being equal, it is rational in this sense to moderate one's degrees of credence (closer to 0.5) in caution-oriented cases and have more extreme credences (closer to 0 or 1) in rush-oriented cases. According to the version of NCP I favour, our natural psychological tendencies provide standards of normalcy approximating the way in which rational credences should be sensitive to practical factors. We could allow some minor deviations from such standards in order to avoid forms of reasoning that we consider irrational or biased. Standards for

¹⁰ The above remarks also help to address a common worry affecting unconstrained versions of pragmatic encroachment according to which such views would lead to a vicious regress (Fumerton, 2010). The regress is avoided if, as explained above, we can be (and often are) sensitive to practical factors even though we do not have doxastic attitudes about such factors, or about practical factors affecting these attitudes. Thanks to an anonymous referee for pressing me to consider this worry.

rational credence would be the result of a reflective equilibrium between psychological findings, intuitive judgments about which cases normally count as rational, and independently motivated constraints (e.g. coherence requirements).¹¹

Consider two notions of rationality: Ideal epistemic rationality and bounded rationality. *Ideal epistemic rationality* takes truth, accuracy and (possibly) evidence as its only standards, abstracting away from limitations of agents' cognitive abilities. In contemporary epistemology, ideal rationality is often taken to be the default notion of epistemic rationality. Could NCP be considered a theory of ideal rationality? Many have argued that this cannot be the case, since CPS makes one vulnerable to diachronic Dutch Books (Clarke, 2013; Greco, 2013; Rubin, 2015; Schroeder, 2018; Shipley, 2021). An agent with practically sensitive credences would accept a certain series of bets, provided respectively before and after certain changes in the relevant practical factors, that collectively lead to a sure loss. For example, suppose that in a low-stakes situation Hannah's credence in p is 0.8. Hannah would be willing to pay \$4 for a ticket that pays back \$5 if p , but decline to pay \$2 for a ticket that pays back \$5 if not- p . Hannah indeed bets \$4 that p . Then Hannah is offered a ticket that costs \$4000 and pays \$5000 if p . Suppose that, due to the influence of perceived high stakes, Hannah's credence in p drops to 0.6 and credence in not- p rises to 0.4. With this new credence distribution, Hannah would now happily accept the previous bet on not- p and pay \$2 for a win of \$5. In the end, Hannah finds herself spending \$6 on a pair of tickets that can only possibly pay back \$5.

According to classical Dutch Book arguments, if a subject is vulnerable to a Dutch Book, then her credences cannot be ideally or maximally rational. Moreover, if Hannah is reflective enough to realise that she is Dutch Book-able in this way, it is hardly coherent for her to accept the first bet on p while also expecting that soon she will consider fair the bet on not- p she initially rejected. Thus CPS cannot be ideally rational, at least if we assume a standard Bayesian picture of rationality.

Even though CPS fails to meet ideal standards of rationality, it can still be considered rational in some other sense. Gao (2016, 2019) argues that CPS manifests *bounded rationality* – roughly, a kind of rationality that takes into account the subject's cognitive limitations and the environmental circumstances.¹² According to Gao, CPS is boundedly rational to the extent that it allows us to allocate our time and

¹¹ For a formal model providing a helpful framework for this version of NCP, see Pettigrew (2016)

¹² A theory of bounded rationality focuses both on the structure of the environments and on the adaptation of the capacities of cognitive systems to the environments. This conception of bounded rationality is developed in studies on 'ecological rationality'. See Todd and Brighton (2016) for a recent development of the theory of ecological rationality and relevant references. For a connection between bounded rationality and self-deception, see Gao (2021).

energy in accordance with the importance and demands of each different task and to form a firm opinion at the point in which the accuracy of judgment is sufficient for a given purpose.¹³

Once we recognize different rationality standards delivering incompatible verdicts, a further open question is what credences one should adopt *all things considered*. Even though we are ready to accept that CPS is convenient and effective in serving some of our practical purposes in a range of circumstances, this is not yet enough to vindicate CPS as the most rational option overall. After all, as Dutch Book arguments show, CPS can also lead to inaccurate estimates and monetary losses. If we want to draw a conclusion about whether CPS is rational *all things considered*, in a substantive normative sense, we must find the way to weigh considerations which involve different types of value (both epistemic and practical) against each other and come out with an overall verdict.

In the next section, I will consider a previous proposal that was precisely meant to address this type of issue: Bishop and Trout's (2005) cost-benefit analysis of reasoning strategies. I will rely on their framework as a prototype for constructing a new framework suitable for assessing the overall rationality of CPS.

3. COST-BENEFIT ANALYSIS AND ITS APPLICATION TO CPS AND CPI

Cost-benefit analysis is based on a very plausible idea: Advantages should be weighed against disadvantages, costs against benefits. A cost-benefit analysis estimates the trade-offs of available options and permits the agent to identify the option that promises the greatest total benefits.

This analysis has also been applied in epistemology. In particular, epistemic consequentialists implicitly assume a cost-benefit analysis model in their arguments.¹⁴ Another application in epistemology, especially relevant for my project, is Bishop and Trout (2005)'s cost-benefit analysis assessing the merits of reasoning strategies. Bishop and Trout's framework is supposed to provide guidance on how people can improve their reasoning practices through adopting new reasoning strategies. In some respects, my approach could be considered a sort of extension of Bishop and Trout's, to the extent that the results of my analysis help to determine whether it is recommendable for our credences to be sensitive to practical factors, how much evidence we should collect and which evidence-weighting strategies we should adopt in order to settle on a given degree of confidence. Given the similarities with Bishop and Trout's project, their cost-benefit framework naturally provides a useful template for my analysis.

Bishop and Trout identify three factors that tend to contribute to the quality of a reasoning strategy: The strategy's reliability, its tractability, and the significance of the problems the strategy is meant to

¹³ See Gao (2019, §3) for an explanation of why bounded rationality can be classified as a sort of epistemic rationality.

¹⁴ For an overview on epistemic consequentialism see Dunn (2015).

tackle (I shall introduce these factors below). The quality of different strategies depends on how they fare along these three dimensions. Based on these criteria, Bishop and Trout's analysis assesses trade-offs between different reasoning strategies. For instance, although adopting a certain strategy could sometimes lead to more reliable outcomes, sticking to a cheaper and less reliable one could save resources and allow the agent to tackle other significant problems more effectively.¹⁵ The choice of whether to change the reasoning strategy would then rest on the overall weight of costs and benefits emerging from such analysis.

Relying on the same methodology, my analysis is supposed to assess CPS versus CPI, as well as their respective methods of credence's regulation, according to some of the same factors individuated by Bishop and Trout. Let me consider in more detail Bishop and Trout's factors and how they could be relevant to my own analysis:

1. *The strategy's reliability.* It concerns the true versus false ratio of predictions that a specific reasoning strategy or method delivers. Bishop and Trout propose that epistemology should recommend *robust reasoning strategies* that lead to accurate predictions across a wide range of environments.

Bishop and Trout refer exclusively to the reliability of cognitive strategies, but since we are assessing credences, in our analysis we have also to consider another related, widely-recognized parameter, namely the *credences' accuracy*, which measures the distance of a credence state or function from truth or perfect accuracy.¹⁶ This captures the idea that the higher the credence we assign to a true (or false) proposition, the better (or worse) it is, which is analogous to the reliabilist idea that the higher ratio of beliefs in truths (or falsehoods), the better (or worse) it is.

2. *The strategy's tractability.* Tractability is about how cognitively demanding it is to employ a certain reasoning strategy. Reasoning strategies that consume fewer cognitive resources and are easier to employ in tackling reasoning problems are better than more demanding and less tractable ones. Switching to a new reasoning strategy often comes at significant costs in terms of time and energy that could have been used for other tasks. Such costs often come in two kinds: (i) *Start-up costs* associated with adopting new strategies, that is, time and energy used for searching and learning to implement a new strategy; (ii) *execution costs* associated with executing a reasoning strategy.

¹⁵ A specific example: the consider-the-opposite strategy involves explicitly considering reasons why one's judgment might be wrong. Adopting this strategy is helpful for containing overconfidence (Plous, 1993, p. 228). But an abuse or over-extensive use of it to solve trivial problems such as whether to change the roll of toilet paper would turn one into an unhappy, neurotic nebbish (Bishop & Trout, 2005, pp. 69-70). Such reasoning patterns would prevent one from pursuing other more meaningful epistemic goals. In this case, reliability (in particular w.r.t. non-significant problems) should be sacrificed for general wellbeing and overall epistemic excellence.

¹⁶ The most common method for measuring the credences' degree of accuracy uses the Brier score.

Start-up costs tend to be conservative epistemic forces that discourage the adoption of new reasoning strategies. In this respect, reasoning strategies that we tend to use by default, such as ‘fast and frugal’ heuristics, have a built-in advantage over strategies that are new to the agent. But start-up costs sometimes can be balanced by low execution costs. For example, learning how to use a new data analysis strategy may be quite difficult, but once one learns how to use it, it may be more handy and effective than using the old one. In such cases, the investment for adopting the new strategy becomes worthy for long-term purposes. As we will see later, the above considerations are highly relevant for our assessment of CPS and CPI.

3. *The significance of the problems the strategy is meant to tackle.* When two strategies have different scope of application, the importance of the problems that each strategy is meant to tackle becomes relevant. However, since CPS and CPI share the same scope of application (i.e. credence regulation in general), they cannot differ regarding this dimension of assessment. For this reason I shall set it aside in my analysis.

From the above discussion it emerges that two factors in Bishop and Trout’s framework are also relevant for my analysis: *Reliability* and *tractability*. However, as mentioned above, methods’ reliability should be supplemented by a further standard, that is, the credence’s degree of *accuracy*. In addition, other factors should be taken into account in the analysis. One such factor, anticipated in Section 2, concerns *monetary losses* that victims of Dutch Books would face. Moreover, we should consider also other *practical non-monetary consequences* of the adoption of CPS and CPI methods.

In sum, my analysis aims to comparatively assess the quality of two competing methods of credence regulation: CPS versus CPI. These methods are assessed along the following dimensions:

- (i) Accuracy and reliability
- (ii) Tractability.
- (iii) Monetary losses.
- (iv) Other practical losses.

The *all-things-considered* best method is determined by the overall weight of costs and benefits relative to the above dimensions. In the next section, I will proceed to evaluate and compare the costs and benefits of the two target methods according to the above listed criteria.

Before proceeding further, two methodological issues are worth addressing. First, one may worry whether it is legitimate to apply a cost-benefit analysis to assess CPS versus CPI. Cost-benefit analyses are traditionally applied to *intentional* actions or decisions, whereas CPS and CPI concern the regulation of doxastic attitudes through (possibly) sub-intentional *cognitive* systems. But such worry applies to Bishop

and Trout's project as well. Their targets of evaluation are reasoning systems, which also concern sub-intentional states and processes. Epistemic utility theory is another prominent example of the application of a similar methodology to evaluate doxastic states rather than actions. Thus, I cannot see any serious reason against an extension of a cost-benefit analysis to cognitive states and systems. At the very least, it is the burden of my opponent to provide such reasons.

Second, in practice, in order to compare relevant options, cost-benefit analyses often place monetary value as the commensurable value units measuring the costs and benefits of each option. A familiar objection to cost-benefit analysis is that some values are incommensurable with money. More generally, there seems to be no "neutral" measure for evaluating entities that we conceive as incomparable. This worry seems also to apply to our cost-benefit analysis: It is unclear how to compare apparently incommensurable factors such as reliability, tractability and practical losses. In response, Bishop and Trout (2005) point out that even if analyses bearing on incommensurable values can be problematic, they are far from fruitless. A cost-benefit analysis plagued by the incommensurability problem can still be useful because it can help to clarify the aims and values implicit in our decisions, and thus help us to adjust our priorities in ways that better reflect our values. Moreover, in some cases even though different types of value may be incommensurable, one option may still obviously prevail because its benefits are much more significant than those of other options and its costs are negligible. I will argue that this is precisely the case in my analysis.

4. A COST-BENEFIT ANALYSIS OF CPS VERSUS CPI

In this section, I first analyse the costs and benefits of CPS (Section 4.1) and CPI (Section 4.2) on the basis of the parameters introduced in the previous section. I then compare the outcomes (Section 4.3).

4.1. Costs and benefits of CPS

4.1.1. Accuracy and reliability

It may be argued that updating credences exclusively on new evidence would deliver a higher degree of accuracy. Many have pointed out that CPS implies systematic violations of Bayesian conditionalization principles. This would arguably come with losses in accuracy. In particular, according to Greaves and Wallace (2006), one can maximize expected epistemic accuracy only if one updates one's credence given conditionalization principles. Furthermore, Briggs and Pettigrew (2020) and Nielsen (2021) argue that violations of conditionalization would make one vulnerable to a sure accuracy loss.¹⁷

¹⁷ In addition to conditionalization, CPS may also imply violations of another alleged constraint on rational credence, the Principal Principle. It is highly contentious whether CPS leads to systematic violations of the principle, and if so whether this leads to serious practical costs. Space constraints prevent me from providing an adequate

However, it is not obvious that CPS is committed to a violation of conditionalization rules and thus to a sure accuracy loss. Elsewhere I have shown how contextual factors can be factored into legitimate Bayesian methods of credence update, preserving coherent credence functions (Gao 2019).¹⁸ In particular, I argue that CPS is fully compatible with updating credences on the basis of conditionalization rules—though the input in conditionalization would not be just evidence but partly a factor of other contextual features. For these reasons, I am not convinced that CPS would systematically deliver less accurate credences compared to a version of CPI updating only on new evidence. Nonetheless, for the sake of argument, in what follows I shall grant this point to my opponents.

Although CPS might fall short of delivering maximally accurate attitudes, it produces a degree of accuracy that is good enough for most practical purposes. In particular, CPS allows us to automatically adjust the strength of evidential support given specific practical demands of the situation. For instance, a tendency to be under-confident when we face critical decisions ensures that we won't rush to take a decision with dangerous and irreversible consequences. In Peanut Sandwich-B, knowing that my guest is seriously allergic to peanuts would make me less confident about whether there is any peanut in his sandwich, thus forcing me to buy it only after a careful double-check.¹⁹ Conversely, the tendency to be overconfident on matters with no potentially serious practical consequences or on urgent matters prevents us from wasting cognitive resources on issues that do not really deserve it. For example, if my assessments of the students' homework only make a very slight difference to their final evaluation, being slightly overconfident and reaching a fast judgment instead of taking a careful and in-depth reading would be a smart and reasonable option.

In sum, the possible accuracy-loss that goes with CPS does not seem to have particularly serious consequences for most practical purposes. Moreover, it does not come for nothing, as it is traded for more robust margins of safety from error in ordinary high-stakes situations (as well as other cases involving caution-oriented factors).²⁰ Conversely, the possible loss in accuracy is traded for

discussion of this issue here. See Gao's *Shifty credence: A defence of credal pragmatism* (unpublished manuscript) for further discussions.

¹⁸ See also Pettigrew (2016) for a probabilistically consistent version of CPS.

¹⁹ Couldn't we instead explain stakes-sensitivity to decision-making in alternative ways, for instance by means of risk-averse decision theories representing stakes as decreasing marginal utilities? From the perspective of a cost-benefit analysis, it is far from obvious that the use of decision theory would constitute a better way of making our decisions sensitive to stakes. Decision theory is a highly sophisticated technical tool. Its implementation in ordinary decision-making would have much higher costs compared to standard decision strategies used by lay people in normal circumstances. More conventional and natural decision methods (e.g., relying on heuristics and our automatic built-in credal sensitivity to practical factors) are more cost-effective and have much fewer implementation costs.

²⁰ Assuming a threshold view about outright belief, a lower degree of credence in caution-oriented cases also keeps one safer from false belief. Thus the gain in safety is not just practical, but also epistemic.

overconfidence when we aim to settle on an opinion urgently (as well as other cases involving rush-oriented factors).

It is also worth noting that CPS is often accompanied by changes in our cognitive strategies, such as reliance on simple heuristics instead of controlled deliberation and vice versa, and by imposing higher or lower evidential thresholds for settled judgment.²¹ Thus we should also recognise variable degrees of reliability of these different cognitive strategies that are variably deployed depending on the practical circumstances. In general, in circumstances featuring caution-oriented factors we tend to use methods that are typically more reliable than those we use in circumstances with rush-oriented factors. Again, the relatively low reliability of methods used in the latter cases is traded for speed and efficiency.

To sum up, although CPS may involve losses in the degree of accuracy and reliability, such costs in general do not bring about significantly negative practical consequences. In rush-oriented cases these costs are balanced by important practical benefits. Similarly, in caution-oriented cases CPS tends to protect us from serious mistakes by providing us with more robust margins of safety from error.

4.1.2. Tractability

As mentioned in Section 1, there is empirical evidence that we actually form, maintain and revise degrees of confidence in ways that are sensitive to practical factors (see Gao (2019) for an overview). As a natural default method of credence regulation, CPS does not incur any start-up costs related to the implementation of new cognitive strategies. The execution costs of CPS also seem to be negligible: One does not need to make extra efforts to learn CPS strategies given that these are already built-in in our cognitive system. Thus, CPS is perfectly tractable, since it doesn't incur excessive waste of cognitive resources that could be instead used to solve other problems and pursue more significant tasks.

4.1.3. Monetary loss in Dutch Book cases

As explained in Section 3, CPS implies being vulnerable to diachronic Dutch Books. However, as many have recognised, the step from hypothetical betting losses to actual sure losses is far from straightforward.

A first reason to doubt that the vulnerability to Dutch Books would lead to high costs concerns their familiar presupposition that an agent's credences can be associated with betting quotients. The idea that degrees of belief can be read off a subject's betting dispositions was common amongst early Bayesians such as Ramsey and DeFinetti, who advanced operationalist definitions of credence in terms of betting dispositions. This operationalist tendency reflected the influence of popular views of the time, such as logical positivism and behaviourism. However, these positions have been the target of serious criticisms

²¹ See Nagel (2008, 2010) for summaries of empirical studies about how practical factors play an important role in determining how much cognitive effort (such as collecting evidence and the strategy of evidence-weighting before settling our minds) one is willing to allocate in order to reach a settled opinion.

and have long been abandoned. The idea that actual betting behaviors are a reliable measure of one's credence has been challenged in multiple ways. For one thing, it has been convincingly argued that betting dispositions are heavily affected by factors other than credence, such as risk aversion, motives besides money, and the format of the bet (e.g., Earman, 1992; Weatherson, 1999; Christensen, 2001). Others have argued that the very practice of gambling involves prudential and moral dimensions that inevitably misrepresent our actual degree of confidence (Salas, 2019). In addition, Eriksson and Rabinowicz (2013) convincingly argued that intuitive judgments about betting at most capture the degree of belief in a conditional that a proposition p would be the case if the agent were to bet on p , where the belief in this conditional itself is conditioned on certain stipulations (e.g. the opportunity to bet, how suspicious the bettor is, etc.) that inevitably alter the subject's epistemic state, so that an agent's choices in gambling cannot reflect her original credences.²² While these considerations do not decisively show that one's degrees of credence can never be manifested in one's betting dispositions, they seriously question the idea that ordinary people's actual credences can be associated with the special kind of betting quotients presupposed by Dutch Book arguments.

Second, being Dutch-booked requires confronting a clever and fully informed bookie who could, and would, take advantage of you. Such a qualified bookie should have two specific types of knowledge: (i) Relevant technical knowledge and abilities required for constructing a profitable set of bets; (ii) knowledge about the subject's credence in the relevant propositions before and after certain practical changes occur. Such artificial settings make the threat of clever bookies highly exceptional.²³

Last but not least, from the perspective of the subject being offered such bets, the quiriness and exceptionality of such betting situations would almost certainly discourage the subject from accepting the bet. In addition, given our natural risk-aversion, the more money one could lose by accepting a bet, the more reluctant one would be to accept the bet. Thus, the chances that one actually would find herself in this type of betting scenario and end up losing significant amounts of money seem very low and far-fetched.

In conclusion, being Dutch-booked requires the following three conditions: betting quotients accurately reflecting the agent's degrees of credence, the existence of a qualified bookie, and the agent's disposition to accept similar bets. Each of these conditions is objectively hard to meet in real-world cases. In particular, the prospect of all three conditions being simultaneously satisfied seems ultimately very unlikely and far-fetched. Indeed, straightforward evidence that Dutch Books do not constitute a serious

²² Some have also argued that betting scenarios where the costs of losing are extremely high almost inevitably affect one's evidence for a proposition, thereby modifying our initial epistemic position (and reasonable doxastic attitudes) toward a proposition (Hacking, 1965; Eriksson & Rabinowicz, 2013; Dodd, 2017; Salas, 2019; Fassio, 2020, 2021).

²³ Cf. Vineberg (2016, §1.4). See also Williamson (forthcoming).

and concrete practical threat to CPS comes from a simple empirical observation: As a matter of fact, our credences are sensitive to practical factors, but we rarely incur heavy monetary losses due to Dutch Books.

4.1.4. Money loss in more realistic scenarios

One may reasonably worry that in other more realistic scenarios CPS could lead to monetary losses. One example concerns insurance purchases. Usually, insurance companies sell policies to cover the costs of possible catastrophic events, such as traffic accidents, serious illnesses and damages to high-valued properties. CPS could make us prone to be overconfident that such high-stakes events could happen, assessing them as more probable than they really are. For example, even though the probability of dying in a plane crash is around one in 188,364 (Ng & Ang, 2019), people may still fear the possibility of this event, overestimate its probability and buy a flight safety insurance that costs \$10.

Is it irrational to buy insurance even if it enables insurance companies to financially exploit us? Even though there is room for arguing that such purchases are not optimal, ultimately few people would judge such behaviours as irrational. Insurance provides us essential financial back-up in catastrophic situations. Although insurances may be often overpriced, the amount of money that one pays for them is usually insignificant compared to one's income and cannot make a real difference to one's living quality. On the contrary, if the disastrous situation happened, without insurance the incurred loss would be huge. Furthermore, buying insurance also has psychological benefits. The feeling of security itself has an important positive value for the agent, which is certainly worth some small money loss. Thus, even if credence's practical sensitivity can lead to small money losses in insurance purchases and similar cases, the purchases themselves are quite reasonable and fully affordable.

Consider a second example. In stock investments CPS seems to prevent one from buying when prices are lowest and selling when they are highest. When stock prices continue going down in a recession, one tends to be pessimistic (high perceived-stakes) and overestimate the probability of the downward trend of prices in the near future, whereas when stock prices shoot up in a boom one tends to be optimistic (low perceived-stakes) and underestimate the probability of the downward trend. Thus, CPS prevents one from objectively evaluating the trends and acting with optimal investment timing.

In response, I suggest that we should make a distinction between lay investors and experts. The reasoning strategy adopted by experts is very different from the one used by lay investors. Most lay investors rely on their hunches or confidence in their financial decisions. Credence naturally constitutes an important ground in their reasoning practices. In this regard, CPS actually helps to prevent us from making reckless decisions and gambling on shaky evidential grounds. In particular, thanks to CPS, the more we plan to invest the more cautious we tend to be in our decisions. By contrast, experts are already used to employing more accurate and professional predicting tools. The information they use is constituted by objective parameters rather than subjective conjectures. Hence, credence and its sensitivity

doesn't play any significant role in their decision-making and don't bring about any significant consequence.

In sum, in cases like insurance purchases and lay people's stock investments, in which CPS may be responsible for suboptimal monetary decisions, the incurred monetary loss tends to be either insignificant or otherwise beneficial. Some may contend that there are other realistic situations in which CPS could lead to significant practical losses. While this may be the case, the burden of proof is on my opponents.

4.2. Costs and benefits of CPI

4.2.1. Accuracy and reliability

According to CPI, degrees of credence should straightforwardly correspond to the strength of evidence available to the subject. CPI doesn't lead to the violation of standard Bayesian updating rules, and could deliver a maximal degree of expected accuracy.

For what concerns methods' reliability, one (maybe the only) straightforward way to achieve CPI is by updating one's credence by strictly conditionalizing on new evidence. Since updating credence by evidential conditionalization is supposed to deliver more accurate attitudes, it tends to produce a higher true versus false ratio of outputs. So, in principle, methods implementing CPI are more reliable than those used by CPS. However, as I will argue below (see Tractability), Bayesian methods are hard to apply in practice. Given the great difficulties that human beings have in reasoning with Bayesian rules, implementing Bayesian updating methods will almost certainly fail to lead to an actual increase in reliability. Even worse, using evidential conditionalization rules in everyday practical reasoning could produce less reliable outputs than most CPS methods used by ordinary people. In many ordinary cases, the use of CPS methods will provide a higher degree of reliability compared to CPI ones.

Moreover, as mentioned in Section 4.1, CPS guarantees a wider margin of safety from error in high-stakes situations (and other caution-oriented cases). In contrast, in these situations forming credence by conditionalizing exclusively on evidence does not guarantee a margin of safety from error as wide as CPS methods. Given how unreliable the implementation of CPI methods could be in practice, this may even result in a low degree of safety, which may lead to imprudent decisions and seriously bad consequences.

4.2.2. Tractability

The implementation of CPI may seem pretty straightforward: One should just systematically deploy invariant cognitive strategies exclusively sensitive to evidence. However, things are not so simple. Our default methods of credence regulation are sensitive to practical factors. CPI can only be achieved by deliberately implementing new cognitive strategies.

In particular, in order to screen off the influence of practical factors in our doxastic regulation, one has to be self-vigilant about the presence of the relevant practical factors and supervise the formation,

revision and retention of one's credences when such factors are present. The relevant practical factors are several. They include, amongst others, stakes, the availability of further evidence, time pressure, dullness of the task and environmental noise. Moreover, some of these factors can often be perceived only at a subconscious level. Keeping track of all these factors, in particular inconspicuous ones, isn't a simple task for any actual agent. This raises the execution costs of CPI.

In addition, once one has identified such factors, one should find a way to screen off their possible effects. But it is hard to see how this could be done. One possible method is to reflectively update credence by conditionalizing on one's evidence. But in most cases we can hardly update our confidence by *explicitly* conditionalizing on our prior credences.²⁴ This is because our evidence (in particular non-statistical information) rarely provides precise probabilistic support, and we are hardly aware of the degrees of our prior credences.²⁵ Even in cases in which prior credences are explicitly accessible or determinable, most humans do not possess the competences for reasoning with probabilities and executing Bayesian conditionalizations. Empirical evidence shows that even intelligent groups of people (including persons with a proper training in probability calculus) systematically make probabilistic mistakes (e.g., Casscells et al., 1978; Spiegelhalter, 2019, p. 209). Given the above difficulties, we would be more prone to error if we were deliberately following Bayesian updating methods than if we relied on our everyday CPS methods. Thus, the loss in tractability for adopting CPI is not even compensated by a gain in reliability.

One could instead try to screen-off the influence of practical factors in forming and updating one's credences by adopting specific debiasing strategies similar to those proposed to correct other prominent biases. For instance, we could try to become aware of the effects of CPS and reduce such effects through intentional control.²⁶ But in practice these strategies don't seem easily implementable either. No endeavours have been devoted so far to figuring out whether there are any specific debiasing strategies that could effectively screen off CPS effects. Moreover, even if such strategies were found, implementing them would likely be inconvenient and time-consuming.^{27,28}

²⁴ Although Bayesian cognitive science suggests that our perceptual systems are able to execute unconscious inferences conforming to Bayesian norms or approximate Bayesian inferences in a range of tasks (see Rescorla, 2021), it is unrealistic to assume that all or most of our mental processes conform to such norms. See also Williamson (forthcoming). Note also that experimental scenarios where humans execute or approximate Bayesian inference are typically ones that do not feature practical factors that are likely to influence one's credence.

²⁵ Horgan (2017) and Byrne (2022) argue that qualitative degrees of confidence expressed in utterances like "*p* is highly likely" and "*p* is moderately likely" and comparative judgments of confidence instantiated in "*p* is more likely than *q*" and "I think *p* is twice as likely to be true than *q*" fall short of justifying the rich probabilistic structure of quantitative degrees of confidence presupposed by Bayesian epistemology.

²⁶ See Mendoza et al. (2010) and Morewedge et al. (2015) for promising debiasing strategies.

²⁷ Such strategies would be particularly demanding given the variety of relevant practical factors.

²⁸ Credence's automatic regulation may license instances of irrationality such as biases. So CPS could face additional start-up costs required for adopting bias-avoiding strategies. Does not CPS also have tractability costs

In sum, there seems not yet to be a way of implementing CPI that avoids high tractability costs.

4.2.3. Other practical costs

In circumstances featuring rush-oriented factors (e.g. urgency), CPS allows the agent to reach a high degree of confidence, and eventually to form a settled opinion in faster and cheaper ways. In such circumstances, CPI would yield only a lower degree of credence on the basis of the same evidence, often not sufficient for the subject to form a settled opinion. This could have serious costs, such as failing to take an urgent decision within useful time limits.

Conversely, in circumstances involving caution-oriented factors, CPS requires the agent to have relatively stronger evidence in order to form a high degree of confidence, and possibly a full belief. Factors like high costs of error call for exceptionally prudent behaviours and dispositions. In such cases it seems safer to be highly confident, form a settled opinion and rely on these states in one's practical reasoning only on the basis of very robust evidential grounds. This demand for a robust epistemic position could guarantee a wider margin of safety from error compared to normal circumstances. Suppose, for instance, that the cost of being mistaken about p is your death. Even though in normal circumstances your evidence would have been sufficient to form a settled opinion and rely on p , given the current very high stakes this would be highly imprudent.

As mentioned above (see Section 4.2.1), in similar high-stakes circumstances, CPI could possibly deliver more accurate judgments (since the degree of credence would fit the evidence), but it would not preserve the wide margins of safety from error guaranteed by a prudent state of underconfidence. The costs of implementing CPI would thus be less prudent decisions in highly risky circumstances in which a single mistake could mean irreversible catastrophic consequences. In such cases, it seems more than reasonable to trade some accuracy for a supplementary margin of safety.

4.3. Comparison

Having analysed the costs and benefits of CPS and CPI, we are in a position to compare the two competing strategies of credence regulation. The following table summarizes the results in the previous subsections:

similar to CPI? In response, we could admit that our normal credal sensitivity is mostly reasonable, while also conceding that some of our natural tendencies leading to biases should be corrected and require some form of monitoring. This would admittedly imply some modest start-up and execution costs, though incomparable to those of implementing a completely new model completely divorced from normal ways. In this respect, CPS would still be much less costly compared to CPI, which requires massive revisions of one's cognitive mechanisms and strategies.

	<i>Practical sensitivity (CPS)</i>	<i>Practical insensitivity (CPI)</i>	<i>Costs-benefits comparison</i>
Accuracy and reliability	<ul style="list-style-type: none"> • Possibly less expected accuracy • In practice, CPS methods' implementation tends to be equally or more reliable. 	<ul style="list-style-type: none"> • No losses in expected accuracy • In practice, CPI methods' implementation tends to be less reliable. 	CPI might fare slightly better than CPS in terms of accuracy; but tends to fare worse in terms of reliability.
Tractability	Default strategy, highly tractable	Very demanding, low tractability	CPS is clearly better
Monetary loss	Possible Dutch Books. But negligible costs for most practical purposes.	None	Minimal advantage for CPI, but negligible in realistic settings.
Other practical costs	Negligible or non-existent.	<ul style="list-style-type: none"> • Less practically effective in rush-oriented cases. • Leads to less prudent decisions in caution-oriented cases. 	CPS is clearly better.

Table 1. A comparison of costs and benefits of CPS versus CPI

For what concerns reliability, although CPI is in a position to guarantee a maximal degree of accuracy, the degree of reliability of CPI methods is not specified, as it depends on which specific cognitive strategies are implemented and how. It is also worth noting that CPI methods can result in higher reliability only by implementing new complex and energy consuming cognitive methods. The more complex the method is, the more error-prone humans will be in applying it. Thus, even assuming that CPI methods would in principle be more reliable with proper implementation, in practice they tend to be less reliable than CPS methods we already use in our daily life. So there is not a clear winner on the reliability score. Not to mention that any advantage that CPI could enjoy in terms of reliability is likely to be paid for by a corresponding significant loss in tractability. These supplementary costs would add to an already existing remarkable advantage of CPS over CPI with respect to tractability.

Concerning monetary losses, the costs of CPS are negligible due to the very remote and far-fetched possibility for real agents to be exploited by Dutch Books. For what concerns other possible practical benefits and costs, it is very unlikely that CPS could involve significant losses. On the contrary, CPS allows us to be practically effective in rush-oriented cases and to avoid imprudent behaviours by

guaranteeing a wider margin of safety from error in high-stakes situations and other cases involving caution-oriented factors. In these respects, CPI has clear disadvantages.

In sum, the drawbacks that CPS could have on accuracy, reliability and monetary losses seem insignificant or negligible. By contrast, the costs of CPI in terms of tractability and other practical losses are realistic and significant. Even though some of the costs and benefits in our analysis are not easily commensurable, CPS seems to have an overall significant advantage over CPI. Hence, the result of the cost-benefit analysis is that *all things considered* our credences should remain sensitive to practical factors as they now are.

It is worth noting that the current result does not concern only practical considerations. Even if we consider exclusively epistemic tasks and goals, avoiding the pursuit of practically insensitive cognitive strategies would save a remarkable amount of cognitive resources that could be used for other epistemic goals. Such epistemic rewards could easily outweigh the limited potential advantages of CPI in terms of accuracy and reliability.

5. CONCLUSION

In this paper, I have provided a cost-benefit analysis of credence's practical sensitivity and insensitivity. The upshot of my analysis is that credal pragmatism does not only provide a realistic descriptive account of the connection between credence and practical factors, but also aspires to the role of a normative theory of rational credences' formation and regulation. In conclusion, we should not worry too much about the fact that credence's practical sensitivity is not recommended by standards of ideal rationality, nor should we try to get rid of such sensitivity. This sensitivity is not a curse, but a blessing for limited human beings like us. This result fits well with studies in cognitive psychology suggesting that most of our cognitive mechanisms are already optimal or near-optimal for normal human beings in normal environments (Gigerenzer 2018, forthcoming; Lieder and Griffiths 2020).

One may be tempted to think that the same conclusion is generalisable to credal sensitivity to other non-evidential factors. However, we will be in a position to draw such a general conclusion only after analysing costs and benefits of credal sensitivity to each type of non-evidential factor. These further analyses have to be postponed to another occasion.

REFERENCES

- Ahlstrom-Vij, K., & Jeff D. (Eds). (2018). *Epistemic Consequentialism*. Oxford: Oxford University Press.
Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.

- Armendt, B. (2010). Stakes and belief. *Philosophical Studies*, 147(1), 71–87. doi: 10.1007/s11098-009-9451-1
- Bishop, M, and Trout, J. D. (2005). *Epistemology and the Psychology of Human Judgment*. New York, NY: Oxford University Press.
- Briggs, R.A., & Pettigrew, R. (2020). An accuracy-dominance argument for conditionalization. *Nous*, 54(1), 162-181. doi: 10.1111/nous.12258
- Byrne, A. (2022). Perception and probability. *Philosophy and Phenomenological Research*, 105(2), 343-363. doi: 10.1111/phpr.12768
- Casscells, W., Schoenberger, A. & Graboys, T. B. (1978). Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine*, 299(18), 999–1001. doi: 10.1056/NEJM197811022991808
- Christensen, D. (2001). Preference-based arguments for probabilism. *Philosophy of Science*, 68(3), 356–376. doi: 10.1086/392889
- Clarke, R. (2013). Belief is credence one (in context). *Philosopher's Imprint*, 13(11), 1–18. doi: hdl.handle.net/2027/spo.3521354.0013.011
- Dodd, D. (2017). Belief and certainty. *Synthese*, 194(11), 4597–4621. doi: 10.1007/s11229-016-1163-4
- Dunn, J. (2015). Epistemic consequentialism. In *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/epis-con/>.
- Earman, J. (1992). *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Eriksson, L., & Rabinowicz, W. (2013). The Interference problem for the betting interpretation of degrees of belief. *Synthese*, 190(5), 809–830. doi: 10.1007/s11229-012-0187-7
- Fassio, D. (2020). Moderate skeptical invariantism. *Erkenntnis*, 85(4), 841–870. doi: 10.1007/s10670-018-0053-1
- . 2021. A (partial) defense of moderate skeptical invariantism. In Kyriacou, C., & Wallbridge, K., (Eds.), *Skeptical Invariantism Reconsidered* (pp. 154–171). New York, NY: Routledge.
- Fumerton, R. (2010). Fencing out pragmatic encroachment. *Philosophical Perspectives*, 24, 243–253. doi: 10.1111/j.1520-8583.2010.00192.x
- Gao, J. (2016). *Belief, knowledge and action* (Doctoral dissertation, University of Edinburgh).
- . (2019). Credal Pragmatism. *Philosophical Studies*, 176(6), 1595–1617. doi: 10.1007/s11098-018-1081-z
- . (2021). Self-deception and pragmatic encroachment: A dilemma for epistemic rationality. *Ratio*, 34(1), 20-32. doi: 10.1111/rati.12288
- . (forthcoming). Credal sensitivism: threshold vs. credence-one. *Inquiry*, 1-22. doi: 10.1080/0020174X.2021.1933744

- Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope with Uncertainty*. Oxford: Oxford University Press.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalization: conditionalization maximizes expected epistemic utility. *Mind*, *115*(459), 607–632. doi: 10.1093/mind/fzl607
- Greco, D. (2015). How I learned to stop worrying and love probability one. *Philosophical Perspectives*, *29*(1), 179–201. doi: 10.1111/phpe.12059
- Hacking, I. (1965). *Logic of Statistical Inference*. Cambridge, MA: Cambridge University Press.
- Hajek, A., & Lin, H. (2017). A tale of two epistemologies? *Res Philosophica*, *94*(2), 207–232. doi: 10.11612/resphil.1540
- Horgan, T. (2017). Troubles for Bayesian formal epistemology. *Res Philosophica*, *94*(2), 233–255. doi: 10.11612/resphil.1535
- Jackson, E. (2019). Belief and credence: why the attitude-type matters. *Philosophical Studies*, *176*(9), 2477–2496. doi: 10.1007/s11098-018-1136-1
- Johnson, E., Hershey, J., Meszaros, J., & Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty*, *7*, 35–51. doi: 10.1007/BF01065313
- Kim, B. (2017). Pragmatic encroachment in epistemology. *Philosophy Compass*, *12*(5), e12415. doi: 10.1111/phc3.12415
- Kruglanski, A., & Webster, D. (1991). Group members' reactions to opinion deviates and conformists at varying degrees of proximity to decision deadline and of environmental noise. *Journal of Personality and Social Psychology*, *61*(2), 212–225. doi: 10.1037/0022-3514.61.2.212
- Kruglanski, A., Webster, D. & Klem, A. (1993). Motivated resistance and openness to persuasion in the presence or absence of prior information. *Journal of Personality and Social Psychology*, *65*(5), 861–876. doi: 10.1037/0022-3514.65.5.861
- Lichtenstein, S. (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory*, *4*(6), 551–578. doi:10.1037/0278-7393.4.6.551
- Lieder, F., & Griffiths, T. (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*(e1): 1–60. doi: 10.1017/S0140525X1900061X
- Loewenstein, G., Weber, E., Hsee, C., & Welch, N. (2001). Risk as feeling. *Psychological Bulletin*, *127*(2), 267–286. doi: 10.1037/0033-2909.127.2.267
- Mayselless, O., & Kruglanski, A. (1987). What makes you so sure? effects of epistemic motivations on judgmental confidence. *Organizational Behavior and Human Decision Processes*, *39*(2), 162–183. doi: 10.1016/0749-5978(87)90036-7

- Mendoza, S., Gollwitzer, P., & Amodio, D. (2010). Reducing the expression of implicit stereotypes: reflexive control through implementation intentions. *Personality and Social Psychology Bulletin*, 36(4), 512–523. doi: 10.1177/01461672103627
- Morewedge, C., Yoon, H., Scopelliti, I., Symborski, C., Korris, J., & Kassam, K. (2015). Debiasing decisions: Improved decision making with a single training intervention. *Policy Insights from the Behavioral and Brain Science*, 2(1), 129–140. doi: 10.1177/237273221560886
- Moss, S. (2018a). IX—Moral encroachment, *Proceedings of the Aristotelian Society*, 118(2), 177–205. doi: 10.1093/arisoc/aoy007
- . (2018b). *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279–294. doi: 10.1080/00048400801886397
- Ng, D., & Ang, A. (2019). How you can survive a plane crash (it's not as rare as you think). *CNA Insider*, November 3, 2019. <https://www.channelnewsasia.com/news/cnainsider/how-you-can-survive-a-plane-crash-it-s-not-as-rare-as-you-think-12058000>.
- Nielsen, M. (2021). Accuracy-dominance and conditionalization. *Philosophical Studies*, 178(10), 3217–3236. doi: 10.1007/s11098-020-01598-6
- Norby, A. (2015). Uncertainty without all the doubt. *Mind and Language*, 30(1), 70–94. doi: 10.1111/mila.12072
- Pettigrew, R. (2016). Jamesian epistemology formalised: an explication of "the will to believe." *Episteme*, 13(3), 253–268. doi: 10.1017/epi.2015.44
- Plous, S. (1993). *The Psychology of Judgment and Decision-Making*. New York, NY: McGraw-Hill.
- Preston, M. G., & Baratta, P. (1948). An experimental study of the auction-value of an uncertain outcome. *The American Journal of Psychology*, 61(2), 183–193. doi: 10.2307/1416964
- Rescorla, M. (2021). Bayesian modeling of the mind: from norms to neurons. *Wiley Interdisciplinary Reviews: Cognitive Science*, 12(1): e1540. doi: 10.1002/wcs.1540
- Rottenstreich, Y., & Hsee, C. H. (2001). Money, kisses, and electric shocks: on the affective psychology of risk. *Psychological Science*, 12(3), 185–190. doi: 10.1111/1467-9280.00334
- Rottenstreich, Y., & Tversky, A. (1997). Unpacking, repacking, and anchoring. advances in support theory. *Psychological Review*, 104(2): 406–415. doi: 10.1037/0033-295X.104.2.406
- Rubin, K. (2015). Total pragmatic encroachment and epistemic permissiveness. *Pacific Philosophical Quarterly*, 96, 12–38. doi: 10.1111/papq.12060
- Rysiew, P. (2008). Rationality disputes—psychology and epistemology. *Philosophy Compass*, 3(6), 1–24. doi:10.1111/j.1747-9991.2008.00178.x
- Salas, J. (2019). Extreme betting. *Ratio*, 32(1), 32–41. doi: 10.1111/rati.12217

- Schroeder, M. (2018). Rational stability under pragmatic encroachment. *Episteme*, 15(3), 297–312. doi: 10.1017/epi.2018.24
- Shipley, J. (2021). Thick credence and pragmatic encroachment. *Philosophical Studies*, 178(2), 339–361. doi: 10.1007/s11098-020-01434-x
- Slovic, P., Fischhoff, B., & Lichtenstein, S. (1982). Facts versus fears. In Kahneman, D., Slovic, P., & Tversky, A., (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*, (pp. 463–492). Cambridge: Cambridge University Press.
- Spiegelhalter, D. (2019). *The Art of Statistics: Learning from Data*. London: Penguin UK.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanovich, K. E. (1999). *Who Is Rational? Studies of Individual Differences in Reasoning*. Mahwah, NJ: Erlbaum.
- Stein, E. (1996). *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*. New York, NY: Oxford University Press.
- Todd, P., & Brighton, H. (2016). Building the theory of ecological rationality. *Minds & Machines*, 26, 9–30. doi: 10.1007/s11023-015-9371-0
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131. doi: 10.1126/science.185.4157.1124
- . (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90 (4), 293–315. doi: 10.1037/0033-295X.90.4.293
- Tversky, A., & Koehler, D. J. (1994). Support theory. a nonextensional representation of subjective probability. *Psychological Review*, 101(4), 547–567. doi: 10.1037/0033-295X.101.4.547
- Vineberg, S. (2016). Dutch Book arguments. In Zalta, E. N., & Nodelman, U., (Eds.), *The Stanford Encyclopedia of Philosophy (Spring 2016 Edition)*. URL = <<https://plato.stanford.edu/archives/fall2022/entries/dutch-book/>>.
- Weatherson, B. (1999). Begging the question and Bayesianism. *Studies in History and Philosophy of Science*, 30, 687–697. doi: 10.1016/S0039-3681(99)00020-5
- Webster, D. (1993). Motivated augmentation and reduction of the overattribution bias. *Journal of Personality and Social Psychology*, 65(2), 261–271. doi: 10.1037/0022-3514.65.2.261
- Williamson, T. (forthcoming). Knowledge, credence, and the strength of belief. In Flowerree, A., & Reed, B., (Eds.), *Expansive Epistemology: Norms, Action, and the Social World*. London: Routledge.