Suspension as Delay: Toward a Decision-Theoretic Credal Theory

Abstract

There has recently been an active debate about the nature of suspension. Many philosophers have criticized the attempt to reduce suspension to credence. This paper develops a new credal theory of suspension based on a decision-theoretic framework that preserves the spirit of credal views, while avoiding their problems.

Keywords: Suspension, delay, credal thresholds, epistemic utility, decision-theoretic framework.

1 Introduction

There has recently been an active debate about the nature of suspension. Many philosophers have criticized the attempt to reduce suspension to credence. Instead, they have treated suspension as *sui generis* or defined it in terms of *higher-order* attitudes (e.g., Crawford 2004, Friedman 2013a, 2013b, 2017, 2019, Atkins 2017, Masny 2020).

This paper develops a new credal theory of suspension that preserves the spirit of traditional credal views while avoiding their problems. Its central idea is to connect suspension with the action of delay. It develops a delay-based theory of suspension within a decision-theoretic framework and defends a biconditional connection between suspension and delay: one rationally suspends judgment on a proposition if and only if delaying the relevant decision is rational, that is, maximizes expected utility, given one's credence in that proposition.

Section 2 outlines naive credal theories and the five problems they face. Section 3 develops a new credal theory of suspension that connects suspension with delay within a decision-theoretic framework. Section 4 examines how this theory addresses the five problems, and Section 5 concludes.

2 Credal Theories, and Their Problems

Credal theories of suspension aim to explain suspension in terms of credences. They typically assume a tripartite doxastic framework in which belief, disbelief, and suspension are the three fundamental traditional doxastic attitudes. On this framework, belief and disbelief are assumed to stand in the usual negation relation: to believe p is

to disbelieve $\neg p$, and to disbelieve p is to believe $\neg p$. Suspension is taken to be closed under negation: suspending judgment on p requires suspending judgment on $\neg p$. These attitudes then extend to compound propositions by the usual rules for conjunctions and disjunctions.

Within this framework, suspension is treated as a distinct attitude — neither (first-order) belief nor disbelief — and is represented by credences that differ from those corresponding to belief or disbelief.

Following the Lockean thesis, these theories maintain that belief in a proposition p corresponds to sufficiently high confidence in p, disbelief to sufficiently low confidence, and suspension to an intermediate level between the two.[1] Accordingly, they represent belief, disbelief, and suspension as follows (e.g., Hempel 1962, Hájek 1998, Hawthorne 2009, Sturgeon 2010, 2020, Clarke 2013, Greco 2015, Hawthorne, Rothschild, and Spectre 2016, Easwaran 2016, Dorst 2019):

> An agent believes a proposition P iff her credence in P is at least a threshold t (i.e., $cr(P) \geq t$, allowing that $cr(P) = 1$); she disbelieves P iff $cr(P) \leq 1 - t$ (equivalently, $cr(\neg P) \geq t$, allowing that $cr(P) = 0$); and she suspends judgment on P iff $1 - t < cr(P) < t$, excluding 0 and 1.

where the threshold t can be fixed or context-variable.

Although initially attractive, this model faces several well-known challenges and two further difficulties that have received less attention.

The first problem is that the threshold t is often criticized as arbitrary or vague, since naive credal theories provide no principled way to identify its precise value (Stalnaker 1984, Foley 1992, 1993). For example, Smith (2010, 2016) argues that no purely probabilistic threshold can be principled, and concludes that what justifies belief must be something other than probability alone—on his view, normic support.

Many have tried to solve this by offering such a principled way. For example, Hawthorne (2009) develops a qualitative logic of confidence and belief to represent belief with a set of probability functions rather than a single precise one. Hempel (1962) inaugurates this strategy of fixing, via expected epistemic utility, the credal threshold for acceptance. Levi (1967) develops a decision-theoretic account of acceptance, rejection, and suspension shortly after, and Lehrer (1970) one of rules of acceptance; Easwaran (2016) and Dorst (2019) later generalize Hempel's model into the full tripartite threshold structure — belief, disbelief, and suspension. Also, Clarke (2013) and Greco (2015) set $t = 1$, so suspension is represented by credences within the range $(0, 1)$, while belief and disbelief are reduced to credences of 1 and 0, respectively.

But there are further challenges to these credal theories. The second problem is that any theory with a threshold below 1 has problems with accounting for suspension on conjunctions and disjunctions (Stalnaker 1984, Friedman 2013a, 2013b, del Rio 2024).[2]

[1] There are two views about how these credal thresholds should be represented: some favor precise credences, while others favor imprecise ones. For example, Hájek (1998) and Sturgeon (2010) defend imprecise credences as better capturing bounded rationality, whereas Hawthorne, Rothschild, and Spectre (2016) and Greco (2015) argue for precise thresholds. While imprecise credal thresholds are worth exploring, this paper adopts the precise-credence approach.

[2] Christensen (2004) responds to this difficulty by denying that binary belief must be closed under conjunction at all, holding that the constraints of rationality apply instead to degrees of confidence, which must obey the laws of probability. Leitgeb (2014, 2017) responds from within a probabilistic framework, requiring that rational belief correspond to a proposition whose probability remains above a threshold under any

An agent can rationally suspend judgment on several independent propositions, as well as on their conjunctions and disjunctions.

If $t < 1$, suspension on individual propositions does not extend well to their logical combinations. Even if an agent suspends on each of several independent propositions, the probability of their disjunction quickly exceeds t, while the probability of their conjunction approaches 0.

For example, consider two coin flips, each with chance $t - \epsilon$ of landing heads. We should suspend judgment on "the first flip lands heads," since the credence $t - \epsilon$ is below the threshold. The same holds for the second flip. But then, the credence that it will land heads on *either* the first *or* second flip is above t, so the theory counts this as belief. However, intuitively, we can still suspend judgment on this disjunction.

On the other hand, theories that set the threshold at 1 face a different challenge. The third problem is that setting $t = 1$ conflicts with the Principal Principle (del Rio 2024:9-10). The Principal Principle also puts pressure on views that set the threshold at 1. Many chance events have objective probability less than 1, yet it is still rational to believe them — e.g., that the sun will rise tomorrow. A theory requiring suspension whenever the credence falls below 1 would overgenerate skeptical results.[3]

Thus, it seems that no extant credal theory of suspension can solve all three problems at the same time. Beyond these familiar issues, credal theories face two further challenges concerning the nature of suspension itself: its connection to inquiry and its connection to action.

The fourth problem concerns the connection between suspension and **inquiry**. Following the tradition that traces back to Sextus and Descartes, philosophers have claimed that when we suspend judgment on a proposition p, we are disposed to continue inquiring about it (e.g., Friedman 2013a, 2013b, 2017, 2019). However, credal thresholds alone cannot explain why suspension involves such a disposition. Credences only reflect an agent's evidence and priors. One may be disposed to inquire p whether her credence in p is high, low, or intermediate. To explain the connection between suspension and inquiry, a credal theory needs to explain why and how credences can influence the disposition to inquire. However, existing credal theories do not do this.

The fifth problem concerns the connection between suspension and **action**. Intuitively, suspension, like belief and disbelief, provides its own distinctive guide to action. Yet with the notable exception of Levi (1980, 1986), this connection has received little attention in the literature, and it remains unclear what kind of action suspension corresponds to.

We can find inspiration from related discussions about belief. Many philosophers have examined how belief connects to action. For example, Fantl and McGrath (2002, 2009) argue that we are justified in believing a proposition p only if we are in fact rational to act *as if* p. Hawthorne (2004) and Stanley (2005) argue that we are justified in

consistent new evidence—a stability condition that, unlike a threshold on individual propositions, secures closure under conjunction. Stability theory, however, is a theory of belief rather than of suspension, and does not address the connections to inquiry and action discussed below.

[3] In fact, the violation of the Principal Principle reflects a broader concern about violating our intuitions and everyday experience of belief. Belief and disbelief come in degrees. One can rationally hold some beliefs or disbeliefs with stronger confidence than others. For example, belief in p can be rationally weaker than belief in $p \vee \neg p$. This is also an important motivation of adopting a credal theory (Christensen 2004, Titelbaum 2022).

believing p only if we are in fact rational in using p as a *reason* or a premise in our *practical reasoning*. Weatherson (2005: 422) argues that one *believes* p 'iff conditionalizing on p doesn't change any conditional preferences over things that matter'.

These views, known as pragmatic encroachment, suggest that similar rational constraints may also apply to suspension. That is, the rationality of suspending judgment on p may be related to the rationality of actions that are relevant to p. How to explain this relation depends on how we understand the relationship between belief and suspension.

One can rationally act *as if* p even when her credence in p is high, low, or intermediate, since decision-making depends on more than purely epistemic factors. To explain the connection between suspension and action, we therefore need a credal theory that accounts for both the relation between belief and suspension and the relation between decision-making and suspension. In particular, it should explain why the determined credal representations of suspension make the corresponding actions rational. Yet, no existing credal theory meets this need, as none take decision-making — involving both epistemic and pragmatic actions — seriously into account.

In light of these difficulties, a satisfactory credal theory of suspension should preserve the spirit of traditional credal views while avoiding their main difficulties. A decision-theoretic framework offers a natural and necessary basis for achieving this. In the next section, I propose a new credal theory of suspension within such a framework. This theory not only addresses the five problems but also advances our understanding of the nature of suspension.

3 A Delay Theory of Suspension

This section develops a new credal theory of suspension by grounding it in the connection between suspension and the rationality of **delay**.

When making decisions, we often have to decide whether to act or to delay. The more certain we are about an outcome, the more we prefer to act; the less certain we are, the more we prefer to delay and wait for more evidence. Similarly, the more confident we are in a proposition, the more inclined we are to believe it; the less confident we are, the more inclined we are to suspend judgment until more evidence is available. The task, then, is to extend this general idea to credal thresholds of suspension — that is, to develop a credal theory that connects suspension with delay.

The argument proceeds in two parts. First, I explain when delay is rational. Second, I show how suspension relates to delay and what this implies for its credal representation.

Agents frequently encounter decisions where it is rational to pause and gather more evidence before acting.

Consider the following example.

> **The Delayed Bridge Case** Suppose you are standing before a bridge. You can either cross it to claim a prize or go home. Or, you can wait to gather more evidence. If you wait, you can potentially learn whether "the bridge is safe" (p), and then decide between crossing and going home on the basis of the new evidence.

If you act immediately, crossing yields 10 if the bridge is safe and –10 if it is unsafe; going home yields 0 in either case. If you delay, you will learn whether the bridge is safe or unsafe and then act on that information. Since inquiry here is assumed to be conclusive — you will end up either knowing p or knowing $\neg p$ — the probability that you learn p is just your current credence in p, and likewise for $\neg p$. However, delaying incurs a cost of 2 if you eventually choose to cross, since going home yields no utility.

This example can be shown by the following decision tree.

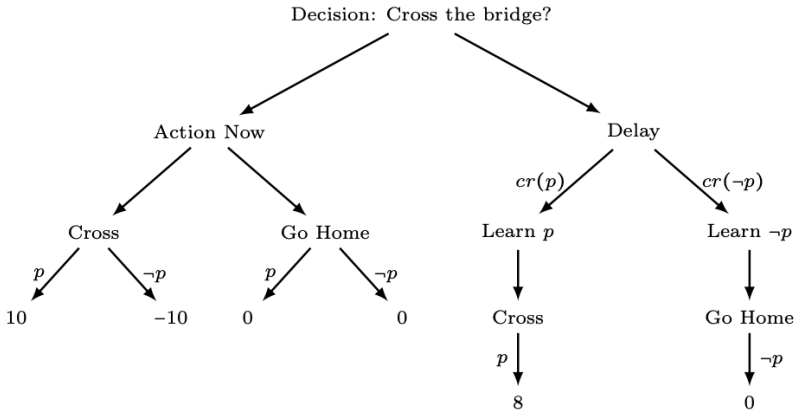

Fig. 1 Decision tree for the bridge case.

Suppose your credence in p is 0.5. Acting immediately then has expected utility 0, whereas delaying yields 4, so delay is rational. If instead your credence in p is 0.9, crossing has expected utility 8, exceeding the utility of delay, so acting immediately becomes rational.

Observation: Depending on the agent's current credence and the payoffs involved, it can be rational to delay.

Specifically, the expected utility of delay is calculated as:

$$EU(delay) = \sum_{i=1}^{k} cr(e_i) \cdot \max_{a \in A} EU(a \mid e_i) - c$$

where $e_1, \ldots, e_k$ are the possible findings of delay (here, that p and that $\neg p$), $cr(e_i)$ is the agent's credence that delay yields e_i (given conclusive inquiry, just her current credence in the corresponding proposition), A is the set of available acts (here, crossing and going home), $EU(a \mid e_i)$ is the expected utility of act a given e_i, and c is the cost of waiting.[4]

[4] More generally, where the cost of delay depends on which act is finally chosen, $EU(delay) = \sum_i cr(e_i) \cdot \max_{a \in A} [\, EU(a \mid e_i) - c(a) \,]$. The bridge case is of this kind: $c(cross) = 2$ and $c(go\ home) = 0$.

The first term is thus the expected value of the information delay provides: the maximum reflects that she acts on what she learns, so the sum measures how much better she can do by choosing in light of her findings rather than choosing now. The second term is what waiting costs her, which may be pragmatic, epistemic, or both.

On this basis, I propose the following theory:

Suspension as Delay (SD) Agent X rationally suspends judgment on a proposition p iff, in a relevant decision problem, her credence in p satisfies

$$EU(\textit{delay to learn } p) \geq EU(\textit{act as if } p \textit{ now})$$

and

$$EU(\textit{delay to learn } p) \geq EU(\textit{act as if } \neg p \textit{ now})$$

except that the endpoints are excluded: when $cr(p) = 1$ or $cr(\neg p) = 1$, X believes or disbelieves p, respectively, rather than suspends judgment.

Applying (SD) to the bridge case, suspension corresponds to the range of credences for which delaying yields higher expected utility than acting immediately. In this example, you rationally suspend judgment when $x \in (0, 0.83]$.[5]

Before explaining the theory further, a few clarifications are needed.

Although the lower threshold in this example is 0, this is not problematic. The thresholds given by (SD) are context-dependent: they reflect the payoffs of the relevant decision problem, rather than general standards for suspension. A threshold of 0 here follows from the particular cost structure of the case: delay costs 2 only if one eventually crosses, so waiting is free on the branch where one ends up going home. At credences near zero, where going home is the likely outcome, delaying therefore costs nothing and remains weakly preferable to acting immediately. Were delay to carry a cost on both branches, the lower threshold would rise above 0.

For a given proposition, the decision problems relevant to determining suspension thresholds may differ across contexts and agents. Consequently, the thresholds may also vary. For a similar view, see Weatherson (2018).[6] I will return to this point in more detail in Section 4.1.

The core of this theory draws on insights from pragmatic encroachment and theories about inquiry but differs from them in important ways. I will discuss these

[5] The thresholds in the bridge case are calculated as follows:

$$\begin{aligned} EU(delay) &\geq EU(cross) \\ 8x &\geq 10x - 10(1-x) \\ x &\leq 0.83 \end{aligned}$$

and,

$$\begin{aligned} EU(delay) &\geq EU(go\ home) \\ 8x &> 0 \\ x &> 0 \end{aligned}$$

[6] There are several nuanced ways of understanding context-dependence, such as context-sensitivity, interest-sensitivity, and subject-sensitive invariantism. In this paper, I will not attempt to distinguish between these interpretations or take a stance on them. Instead, my theory is meant to be compatible with, and flexible enough to work under, most standard versions of these interpretations.

connections and differences in detail in Sections 4.2 and 4.3. In this section, I focus on laying out how to connect suspension with delay and carefully presenting the theory.

With these clarifications in place, I now turn to the theory. The theory assumes a biconditional relationship between suspension and delay as (SD) states. On this assumption, rational suspension can be reduced to the range of credences that make delay rational.

This connection is intuitive, since both suspension and delay involve refraining from taking a decisive stance immediately. To establish the biconditional relationship, let's examine each condition in turn.

On the one hand, delay involves suspension. When one delays the decision, one refrains from making a definitive choice among the available options. In this sense, delay involves a commitment to neutrality about which option is best until the final decision. This neutrality is best explained by suspension rather than belief or disbelief.

Belief and disbelief do not involve any epistemic commitment to neutrality. They take a clear stance on whether a proposition is true or false, and that stance naturally leads to a definite decision. In the bridge example, if you stand for the truth that the bridge is safe, it would be inconsistent to say, 'I believe the bridge is safe, but I will remain neutral and delay my decision,' especially when you know the cost of delay.

By contrast, suspension on a proposition involves refraining from taking a stance between belief and disbelief in it.[7] This commitment to neutrality explains why, when one delays, one also remains neutral among the options. Since the best option depends on the truth of the relevant proposition, the neutrality involved in delay rests on an epistemic commitment to neutrality about whether the proposition is true. As in the bridge case, where the best choice depends on whether the bridge is safe.

On the other hand, suspension involves delay, since the same reasons that justify suspending judgment also justify delaying a decision. Many would agree that one suspends judgment when the available evidence is insufficient to support either belief or disbelief (e.g., Friedman 2013a, 2017, del Rio 2024). Otherwise, continuing to suspend despite already having sufficient evidence is irrational. Saying something like, 'I have sufficient evidence to believe that it will rain tomorrow, but I suspend judgment on it,' sounds inconsistent.

The same reason leads to delay, since deciding with insufficient evidence exposes one to a significant risk of error. Consider the bridge example, if you cross the bridge based on insufficient evidence to support the belief that it is safe, you face a known

[7] This view is widely shared across different accounts of suspension. There are two main understandings: some treat suspension as simply the absence of belief and disbelief, or as believing that one neither believes nor disbelieves (e.g., Salmon 1986, Crawford 2004, Kaplan 2013, Masny 2020), while others treat it as a distinct attitude in its own right, involving a commitment to neutrality and indecision between belief and disbelief (e.g., Sturgeon 2010, Friedman 2013a, 2013b, 2017, 2019, Atkins 2017). Despite their difference, both agree that to suspend judgment on a proposition is neither to believe nor to disbelieve it.

risk of collapse. But if you go home, you will likely miss the prize. So acting now would be irrational; you should delay.[8]

This supports a strong connection between suspension and delay: each is rational when the other is. Theory (SD) is justified by this close relationship. It suggests that thresholds of suspension can vary across contexts, since agents may be rational to delay at different credence levels when facing different decision problems or evaluations. This context-dependence helps address the challenges facing credal theories, and I will explain it further when discussing the last two problems in the next section.

4 Applying the New Theory to Problems for Credal Theories

Having illustrated the new credal theory of suspension, I now turn to examine how it addresses the five central questions for credal theories discussed in Section 2. I highlight its key innovation: it can accommodate the connections between suspension, inquiry, and action.

4.1 Addressing the First Three Problems

Let's begin with the first three problems. I argue that the delay theory of suspension can respond to these challenges by restricting each credal account to a specific context.

First, consider the problem of vague and arbitrary thresholds. The delay framework offers a natural way to avoid it. On this account, the credal thresholds of suspension are the range of credences that make delay rational: when one suspends judgment on p, the thresholds are fixed by the credences on which delay is the expected-utility-maximizing choice in a decision problem that depends on the truth of p.

It is worth seeing why these credences form an interval, and why it falls where it does. Delay is valuable only if what one might learn could change what one does.

[8] One might worry that there are cases of suspension without delay, and delay without suspension. For example, suppose a demon threatens to kill you unless you immediately choose either to cross the bridge or go home, despite your lack of evidence. In such a case, the cost of delay is so high that you must act even without sufficient evidence. Intuitively, the demon's threat should not undermine your rational suspension about whether the bridge is safe: you may still suspend judgment while being forced to act.
On my theory, however, such cases cannot be modeled as instances of rational suspension. If delay is necessary to the neutral and indecisive nature of suspension, then the inability to delay undermines the possibility of genuinely suspending judgment. Since my theory holds that suspension, like delay, is both epistemically and pragmatically context-dependent, the absence of any rational possibility of delay also removes the possibility of rational suspension in that context.
This response may not satisfy everyone. Some views of suspension seem able to handle this problem more easily. For example, higher-order accounts of suspension can explain why one rationally suspends judgment without delaying by appealing to failures of self-knowledge: one may rationally suspend while failing to believe that one is suspending. Naive credal theories of suspension can also accommodate this case, since they do not require suspension to be connected with action.
Some views, however, could struggle here. Views that treat suspension as a *sui generis* attitude connected with inquiry (e.g., Friedman 2013a, 2013b) face a difficulty: they must explain how suspension remains rational when no genuine inquiry is possible. Thus, if this is a problem, it is not unique to my view, but also arises for Friedman's view and for credal theories that aim to explain the connection between suspension and action.
Moreover, the worry can be turned against views that separate suspension from delay. A "revenge demon" of the Dutch Book or Czech Book kind can be constructed: a manipulator who exploits the gap between delay and settled belief by arranging bets that guarantee a loss whenever one delays without suspending judgment. Such cases create pressure specifically for theories that treat suspension and delay as independent. One may also resist the worry by denying that demon-manipulated decision problems are relevant for determining rational thresholds for suspension.

At credences extreme enough that the same act would be best whatever the agent learned, delay yields no informational gain while the cost of waiting remains, so acting immediately would be rational. The interval of suspension is therefore the range of credences at which the expected gain from possibly revising one's choice exceeds the cost of waiting to make it. The thresholds are thus derived rather than stipulated: they lie where these two quantities balance.

This also explains the role played by the cost of delay. By Good's (1967) theorem, the gain from waiting, setting its cost aside, is never negative: costless information cannot leave an agent worse off in expectation.[9] Were delay costless, it would therefore be at least as good as acting immediately at every credence, and (SD) would coincide with the view of Clarke (2013) and Greco (2015) on which belief requires credence 1. The framework thus contains that view as the limiting case in which waiting is free: the more costly waiting is, the narrower the interval of suspension becomes.[10]

These thresholds thus both reflect and explain the biconditional relationship between suspension and delay: the boundary between suspension and belief or disbelief is just the boundary between delaying and acting immediately. They are determined in a principled way that answers to the nature of suspension, and so are neither arbitrary nor vague.

Then, consider the second and third problems — the conjunction/disjunction problem and the violation of the Principal Principle. As mentioned in Section 2, the second and third problems are closely connected. Any theory that addresses the second problem must face the third problem.

The decision-theoretic credal theory avoids both problems because, under (SD), thresholds are defined within particular decision problems, and the resulting suspension intervals are therefore context-dependent. Different propositions can be embedded in different decision contexts, with different stakes or payoffs, so their corresponding thresholds for suspension need not align.

For example, the thresholds for suspension on "the bridge A (a real bridge) is safe" and on "the bridge B (a toy bridge) is safe" may differ, since the two propositions are related to different decision contexts, outcomes, and evaluations. As a result, they lead to different expected utilities for delaying.

In the Bridge A case, acting too soon risks a large loss (–10) if the bridge is unsafe, while delaying incurs a cost of 2. In the Bridge B case, because the bridge is merely a low toy bridge, the stakes are smaller — say, a loss of 4 if the bridge is unsafe and a gain of 4 if it is safe, and the cost of delay is 2. In this case, the expected-utility gap between delaying and acting is therefore smaller. Under these payoffs, the suspension thresholds narrow: the theory gives a credal interval of $(0, 0.67]$ for suspension, which is tighter than the suspension interval in the Bridge A case $(0, 0.83]$.[11]

Within this framework, it does not follow that, simply because credences in conjunctions or disjunctions are straightforward functions of the credences in their

[9] The theorem assumes an agent who updates by conditionalization and maximizes expected utility. It can fail for non-expected-utility theories; see Buchak (2010).

[10] In the bridge case the cost falls only on the crossing branch, so the interval narrows from above alone.

[11] This result is natural and intuitive. When the stakes are high, making a wrong judgment can lead to a large loss. For example, crossing an unsafe bridge could be deadly. In such cases, it is rational to be more cautious — more willing to suspend judgment and wait for further evidence. By contrast, when the stakes are low, the cost of being wrong is small. So, it is rational to take a definite stance with less hesitation.

independent components, the thresholds for suspension on a conjunction or disjunction must exhibit the same functional relation to the thresholds for their components. Thresholds for suspension depend not only on an agent's credence in a proposition, but also on the context — specifically, on how the agent evaluates the utility of delaying within the decision problem related to that proposition.

As the agent shifts attention from one proposition to another — for example, from individual propositions to their conjunction or disjunction — the relevant decision problem changes. The stakes, available actions, and potential payoffs can differ across these contexts, and so the expected utility of delaying can differ as well. For this reason, the suspension intervals associated with different propositions may vary across contexts, and there is no requirement that the interval for the conjunction be wider, narrower, or otherwise systematically related to those for its components.

More generally, there is no principled reason to expect that these thresholds should be collapsed into a single, catch-all standard for suspension. Each threshold is determined by the particular decision problem in which it arises, and different evaluative standards across contexts naturally yield different thresholds In some settings the suspension interval may be (0,1), while in others it may be narrower or stricter. It is also natural for the standards for belief and disbelief to tighten in some contexts — for example, when no objective chance is involved, one might require credence 1 for belief and 0 for disbelief.

At the same time, the framework can accommodate cases in which belief with credence less than 1 (or disbelief with credence greater than 0) is rational. This occurs when, given the relevant decision problem, acting immediately on such a credence maximizes expected utility relative to delaying for further evidence.

Thus, the theory allows for a wide range of rational doxastic responses, each shaped by the particular decision context in which the agent is situated. By permitting extreme thresholds only in the contexts that generate them, the decision-theoretic framework both addresses the conjunction/disjunction problem and avoids the kind of threshold generalization that would conflict with the Principal Principle.

4.2 The Fourth Problem: Suspension and Inquiry

Having addressed the three problems most often discussed in credal theories, we now turn to the fourth and fifth. This subsection considers the fourth.

As noted in Section 2, the fourth problem concerns the connection between suspension and inquiry. Many philosophers hold that suspension essentially involves a disposition to inquire and keep a question open.

Friedman (2013a,b, 2017, 2019) systematically defends that one rationally suspends judgment iff one is *disposed* to further inquiry. She understands inquiry as an attitude directed toward a question, such as curiosity, deliberation, or wondering. It need not involve any active investigation to gather further evidence; it requires only a disposition to keep the question open. Like inquiry, suspension involves a disposition to engage with the *question* and to maintain neutrality about its answer.

The delay theory of suspension is well suited to capture this relationship. On this view, inquiry can be modeled as a special case of delay: an agent inquires when waiting

for additional evidence has higher expected *epistemic utility* than forming a doxastic attitude immediately.

Delay preserves precisely the kind of openness that inquiry requires. When an agent delays forming a belief, she is not merely waiting, but also refrains from settling her view and leaves room for additional evidence to guide her later judgment. How long one delays depends on how long one remains open to receiving new evidence. This "holding open" of the question naturally parallels the interrogative stance of inquiry. Also, like inquiry, delay is not itself the active search for evidence, but it creates the space in which such search and the disposition to the search is possible.

Without delay, it is hard to see how one could remain genuinely open to the question. For example, suppose you claim to suspend judgment on whether the bridge is safe, yet you choose to cross rather than delay. By making that choice, you effectively assume the bridge is safe. Once the decision is made, your attitude toward the question is settled, and the need for further inquiry in that context comes to an end.

As Friedman notes, 'When we suspend some event or activity we put a stop to it, but only temporarily — we put the activity in abeyance, and leave it to be taken up again at a later time'(2017: 16). She observes that suspension involves delay, and that delay carries with it 'some sort of commitment to continued efforts to judge'. This observation fits smoothly with the delay framework: suspension involves temporarily withholding judgment, and delay provides the structural conditions under which further inquiry can occur.

Thus, delay explains why suspension is naturally accompanied by a disposition to inquire: whenever delay is rational, maintaining openness to further evidence is likewise rational.[12]

Let's now examine how the connection between suspension and inquiry can be modeled within decision problems about whether to inquire.

---

[12]In this paper, I model suspension as an attitude toward propositions (e.g., Atkins 2017). On this view, suspension on a proposition p is represented by one's credence in p. For example, suspension on "the bridge is safe" can be represented by a credal interval for that proposition.
Friedman (2013a, b, 2019) however, argues that suspension is directed toward questions rather than propositions. On this view, modeling suspension is less straightforward than in the propositional case. How, if at all, could an agent rationally assign a credence to a question? It seems odd to say, "I assign a credence of 0.5 to whether it is raining today" — what could that even mean?
I propose modeling suspension toward questions in the same way as suspension on propositions. A question can be represented as a set of possible answers, each corresponding to a proposition. The context in which the question is asked determines which answers are relevant, and only those compatible with the context should be included (Stalnaker 1974, Dorst and Mandelkern 2022). For example, consider the question "What will the weather be like tomorrow?" In summer, the relevant answer set may include the propositions "It will rain," "It will be sunny," and "It will be windy." In winter, it may also include "It will snow."
Thus, when an agent suspends judgment on a question, she is suspending judgment on the set of possible answers to that question. Let the answers be propositions $(A_1, A_2, \ldots, A_n)$, where each A_i represents a relevant answer and $n < \infty$ (as the account concerns boundedly rational agents). Suspension on each answer is represented by its own credal interval. Since these answers are in the same context (if they were not, the credal representations could not be aggregated), the credal representation for suspension on the question is obtained by aggregating those for all answers in the set. For example, suppose the question "What will the weather be tomorrow?" has two relevant answers: "it will rain tomorrow" and "it will snow tomorrow." If the credal interval for suspension on the former proposition is $[0.4, 0.6]$, and that for the latter is $[0.3, 0.6]$, then, by taking the lower bound as the minimum of individual lower bounds, the credal interval for suspension on the question is $[0.3, 0.6]$.

In decision-theoretic terms, inquiry is rational when its *expected epistemic utility* exceeds that of forming a settled opinion now. This frames inquiry as an epistemic analogue of the delay option in ordinary decision problems.[13]

In light of the connection between suspension and inquiry, when inquiry is the rational choice, suspension is rational as well. Accordingly, suspension corresponds to the range of credences for which inquiry is the rational option — that is, the credences for which waiting to learn more about p yields higher expected epistemic utility than immediately believing or disbelieving it.[14]

A simplified case illustrates this point. Suppose a scientist investigating the Big Bang hypothesis about the origin of the universe (H). Suppose that, given her current evidence, she faces a decision problem about whether to inquire further into H.

If she chooses to inquire and wait for further evidence, she suspends judgment; whereas if she chooses not to inquire, she effectively settles her opinion by believing or disbelieving H.[15] Suppose she assigns epistemic utilities as follows:

> She receives utility 1 for true belief (i.e., believing H when H is true, or disbelieving H when $\neg H$ is true) — the reward for getting things right;
> she receives utility -1 for false belief (i.e., believing H when H is false, or disbelieving H when $\neg H$ is false) — the penalty for getting things wrong;
> if she chooses to inquire and wait, she incurs a delay cost of 0.2 each time she chooses to delay her decision.

It is worth noting that the epistemic cost of delay differs in an important way from the pragmatic cost of delay. The epistemic cost reflects not only the cost of choosing later, but also the cost of not holding a true belief for a longer time. This is because the time of holding a true belief has its epistemic value. Epistemic utility is not a one-off payoff like claiming a prize. Rather, its effects persist over time: the beliefs one holds at one time shape one's epistemic system at later times and enables more reliable and accurate future updates.

Therefore, holding a true belief for longer yields more epistemic utility. The same holds for false belief: holding a false belief for a longer time is worse than holding it only at one time.

This model incorporates the timing of the agent's decisions. At time t_1, she decides whether to inquire. At t_2, if she did not inquire, she must decide whether to believe or disbelieve based on her current evidence; if she did inquire, she waits to learn new evidence. At t_3, if she did not inquire, she simply continues with the belief or disbelief

[13] The epistemic utility is typically understood in veritist terms, often identified with accuracy. But other interpretations are available, for example, treating it as the quality of an agent's guesses (Horowitz 2017) or as the value of information (e.g., Skyrms 2010). Although the interpretation of epistemic utility is an important issue, it lies beyond the scope of this paper. The decision-theoretic approach developed here neither specifies nor depends on any particular interpretation of epistemic utility. Accordingly, the theory can be compatible with most standard conceptions of epistemic utility.

[14] Being rational in inquiry does not require actively investigating. It requires only a preference for inquiry — a willingness and openness to inquire. I understand this as an interrogative attitude, a disposition toward inquiry.

[15] Based on her theory about suspension and inquiry, Friedman (2019) argues that belief and disbelief are a kind of settled opinion, in the sense that 'a rational believer doesn't also deliberate; she's not curious about some relevant range of questions or wondering about them or contemplating them, and so on' (2019: 312).

she formed at t_2; if she did inquire and learned whether H, she forms a belief or disbelief on that basis.[16]

Here, I assume that the overall epistemic payoff from holding a true belief across time is the *sum* of the payoffs at each time.[17]

The decision problem is represented by the following decision tree, where the labels t_i indicate the time at which the agent makes each decision:

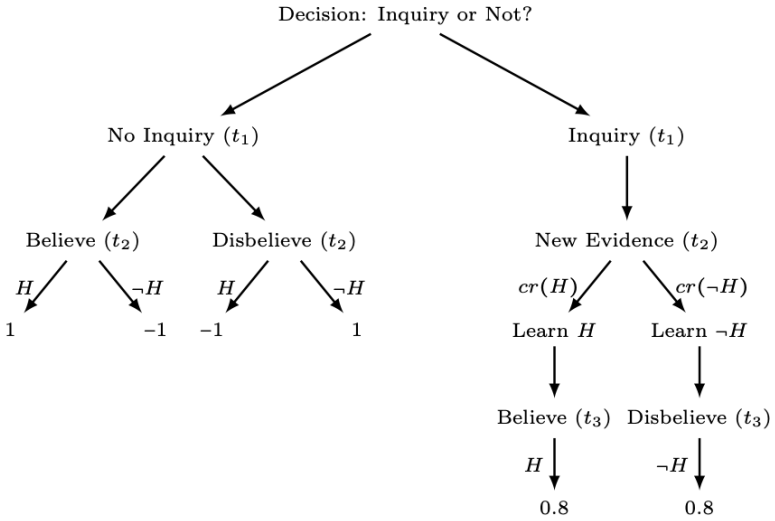

Fig. 2 Decision tree for the scientist's case.

[16]This decision model omits the possibility that the agent learns nothing new by waiting until t_2. In such cases, she must decide again whether to inquire on the basis of her current evidence at t_3. If she chooses not to inquire, she settles her opinion by forming a belief or disbelief; if she chooses to inquire, she continues to wait and incurs an additional cost of delay. The model can accommodate such cases by considering the probability of acquiring new evidence. This probability can be understood either as the objective, truth-oriented chance that the world will reveal relevant information, or as the agent's subjective credence about how likely she is to obtain new evidence.

[17]Some might argue that the overall payoff should be computed non-linearly or according to some alternative rule. This does not threaten the model. The expected utilities here are calculated using one natural set of assumptions, and the framework can readily incorporate different payoff functions if required.

In this context, we can determine the credences corresponding to suspension under theory (SD) by calculating the expected utilities that favor inquiry. Assume the scientist inquires only once. In that case, the suspension interval is $[0.3, 0.7]$.[18]

In this way, the delay theory of suspension offers a natural explanation of the connection between suspension and inquiry: by taking inquiry as a specific case of delay, it is rational to suspend judgment iff it is rational to prefer further inquiry.

My theory can then be stated as follows: you are rationally disposed to inquire about p only if it is rational to delay and wait for new evidence about p. Whether it is rational to delay depends on your credence in p. In a given decision problem that depends on the truth of p, with payoffs held fixed, there is a range of credences that make delay rational — that is, the range at which delay maximizes expected utility. That range determines the credal interval for suspension.

However, it is helpful to note that, although I agree with Friedman in many respects, my theory differs from hers in an important methodological way. Friedman treats suspension as a *sui generis* state. By contrast, my approach treats suspension as something that can be analyzed using decision-theoretic tools. We can model an agent's preferences about inquiry as rational decisions, and thereby connect suspension to the credences that make those preferences rational.

My decision-theoretic approach also differs from standard versions of epistemic utility theory.

First, it captures an important feature that is often overlooked there: the value of *time* in holding a true belief. Because delay is built into the framework, there can be more than one point at which an agent may decide whether to form a belief.

In purely epistemic decision problems, one may form a belief immediately and maintain it over time, or delay and form the belief only later. Holding a true belief for a longer period yields more epistemic utility, since the overall payoff increases each time the belief is maintained. Delay is rational iff its expected epistemic utility exceeds the overall and diachronic expected epistemic utility of forming the belief immediately and holding it throughout.

Second, the framework helps to explain the genuine epistemic value of suspension. Many philosophers use epistemic utility theory to develop credal theories of belief, disbelief, and suspension (e.g., Hempel 1962, Easwaran 2016, Dorst 2019, Pettigrew 2024). However, these theories typically represent suspension as a single credence point that distinguishes belief from disbelief.

---

[18] The upper threshold for suspension is:

$$\begin{aligned} EU(Believe) &\leq EU(Inquiry) \\ 2(2x-1) &\leq 0.8x + 0.8(1-x) \\ x &\leq 0.7 \end{aligned}$$

and the lower threshold is:

$$\begin{aligned} EU(Disbelieve) &\leq EU(Inquiry) \\ 2(-2x+1) &\leq 0.8x + 0.8(1-x) \\ x &\geq 0.3 \end{aligned}$$

.

This way of modeling suspension risks obscuring its distinctive epistemic role. Suspension allows agents to avoid false belief, preserve openness to further evidence, and create time for further inquiry, thereby promoting the long-run pursuit of truth.

My theory of suspension offers a more natural explanation of this value by grounding it in the value of delay. In both pragmatic and purely epistemic cases, the utility of delay (which reflects the epistemic utility of suspension) comes from the possibility of acquiring new evidence and then making more accurate decisions, as well as from the epistemic caution and neutrality that suspension embodies.

In addition, my approach has an advantage over connecting it directly with inquiry: the connection between suspension and delay includes, but goes beyond, the connection between suspension and inquiry.

Delay involves not only a disposition toward inquiry but also a state of pure neutrality. In this neutral state, the agent is indifferent between acting now and delaying, and has no motivation to inquire further. This occurs when the expected utilities of all options are equal and the agent's evidence is sufficiently strong and supports each option equally. In such cases, one can rationally choose to delay purely due to neutrality, which explains why the credal interval of suspension includes the threshold points.

This kind of neutral delay can arise even when one expects that little or no further evidence will become available. For example, this can occur in high-stakes cases. Consider a fundamental scientific theory: if one were to believe it and it later turned out to be false, this could undermine an entire knowledge system. In such high-stakes cases, even when no new evidence is expected and there is no genuine disposition to inquire, it can still be reasonable to delay and suspend judgment, simply to avoid the serious epistemic cost of forming a false belief too quickly.

Moreover, the delay theory can accommodate suspension influenced not only by epistemic factors but also by pragmatic factors. Suspension is not only a matter of searching for evidence; it can also reflect caution and hesitation to believe. For example, when crossing a bridge could cost your life, it can be rational to delay and thus suspend judgment about whether the bridge is safe, even when no new evidence is likely to arise and further inquiry would make no sense. I will explain this in the next section.

Thus, for those who accept the connection between suspension and inquiry, the connection between suspension and delay seems at least as plausible.

4.3 The Fifth Problem: Suspension and Action

The fifth problem concerns the connection between suspension and action.

It is intuitive that each of suspension, belief, and disbelief plays a distinctive role in guiding rational choice. As is often observed, belief, disbelief, and suspension guide us toward different actions. Consider again the bridge example. If you believe that the bridge is safe, it seems natural and reasonable to cross it. If you disbelieve, it seems appropriate to go home. And if you suspend judgment, it seems reasonable to delay your decision.

As discussed in Section 3, the delay theory of suspension provides and explains a principled guide to action when one suspends judgment: one rationally suspends judgment on a proposition p iff one rationally delays making the decision that depends

on the truth of p. The rationality of suspension is affected by the rationality of delay, and vice versa.

This has a precedent in Levi (1980), who also treats doxastic states as bearing directly on rational choice. On his view, the agent's corpus of full beliefs is assigned probability 1 and serves as her standard for serious possibility: the negations of what she fully believes are ruled out in deliberation. When the corpus includes neither h nor $\neg h$, then both are consistent with her corpus and remain seriously possible (1980, pp. 3–7, 37).

Levi and I agree, then, that doxastic change is itself a species of decision-making. We differ, however, in three respects. First, over which utilities fix the thresholds: Levi holds that the utilities governing acceptance are purely epistemic, trading truth against the value of information under an index of caution (1967);[19] (SD) allows the relevant decision problem to be practical as well as epistemic, so that stakes and payoffs also help determine when suspension is rational. Second, over the relation between belief and credence: for Levi, full belief is not a high credence but the corpus against which credences are assigned, and so receives probability 1; (SD) is a credal theory throughout, and can accommodate belief at credences below 1, which better fits the thought that believing something need not amount to being certain of it. Third, over the temporal structure of the explanation: Levi explains suspension synchronically, by the expected epistemic utility of accepting a hypothesis given one's present corpus, whereas (SD) explains it diachronically, by the expected utility of waiting for evidence one does not yet have.

This theory can also fit well with existing theories of pragmatic encroachment on belief, which holds that the rationality of belief can be affected by the rationality of action.[20]

One way to put this point is to treat suspension as precisely the attitude that exclude belief and disbelief, given the assumption of a tripartite doxastic framework. On this understanding, the delay theory of suspension implies that an agent rationally believes or disbelieves p iff one rationally acts immediately, rather than delay, in the relevant decision problem. For example, in the bridge case, belief corresponds to crossing immediately, while disbelief corresponds to going home immediately. This aligns naturally with the existing pragmatic encroachment theses (e.g., Fantl and McGrath 2002, 2009; Weatherson 2005; Hawthorne and Stanley 2008; Ross and Schroeder 2014).

Let us illustrate this with a well-known version of pragmatic encroachment. Fantl and McGrath (2002) argue that Agent X rationally believes a proposition p only if X rationally prefers to act as if p is true. Since disbelieving p can be understood as believing $\neg P$, the same connection applies equally to disbelief and action.

Consider the bridge case. If the proposition p ("the bridge is safe") is true, the rational action is to cross the bridge immediately, since this yields the maximum expected utility of 10. According to Fantl and McGrath's theory, believing p is rational only if it is rational to act as if p is true — in this case, to cross immediately. Thus,

[19] The same holds for contemporary epistemic utility theories such as Easwaran (2016) and Dorst (2019).

[20] The theory also makes room for moral encroachment, the view that moral considerations can affect when suspension is rational. For example, when a belief risks reflecting racist or sexist discrimination, suspension may be required (e.g., Moss 2018; Basu and Schroeder 2019). However, this thesis is more controversial than its pragmatic counterpart, so this paper focuses only on the latter.

belief in p corresponds to credences that make crossing immediately rational. Likewise, disbelief in p corresponds to credences that make going home immediately rational. If suspension is the attitude that excludes both belief and disbelief within a tripartite framework, then suspension should correspond only to the credences that make delay rational. In this way, the delay theory of suspension fits naturally with their pragmatic encroachment thesis.[21]

However, some might wonder whether the delay theory of suspension is just a simple extension of pragmatic encroachment on belief. I argue that it is not, since they are different in important ways.

First, pragmatic encroachment theories do not discuss suspension. By requiring belief to correspond to acting as if p is true and disbelief to acting as if p is false, they leave no distinct room for suspension within their action-based framework. The delay theory of suspension helps articulate that missing connection by showing how suspension can also be modeled within a decision-theoretic structure by connecting it with delay.

Second, the biconditional connection between belief and action implied by the delay theory is not identical to those proposed by pragmatic encroachment theorists, but it involves and helps explain them. Weatherson (2005) defends a biconditional relation between belief and action, though he does not require that the rationality of belief depend on the rationality of action. By contrast, Fantl and McGrath (2002, 2009) and Ross and Schroeder (2014) argue only for a necessary condition on the connection between belief and action.

Third, pragmatic encroachment theories have emphasized that believing p involves treating p as fixed for the purposes of action. By contrast, the delay theory of suspension holds that suspending judgment on p involves *not* treating p as fixed in action.

---

[21] As discussed in Section 2, philosophers defend different versions of this belief–action connection. Besides the version discussed in the main text, two other influential formulations are worth noting. Hawthorne (2004) and Stanley (2005) argue that one rationally believes p only if one rationally prefers to act by treating p as a premise in practical reasoning. Weatherson (2005) argues that one *believes* that p iff one prefers action A to B given that the proposition q is true, and continues to prefer A to B given that $p \vee q$ is true.
Admittedly, the version of the belief–action connection we adopt matters for how we understand pragmatic encroachment on belief. Nevertheless, the delay theory of suspension can fit all three views.
Hawthorne and Stanley hold that a rational believer must treat what she believes as a premise in practical reasoning, though she need not act on it directly. Consider an example: you face two bridges, A and B, and you believe each is safe with the same credence less than 1. If you reason solely as if A is safe, you should cross A. Yet intuitively it is rational to cross either bridge. Their view explains this by allowing beliefs to function as *reasons* for action without being *exclusive* reasons. The fact that A is safe does not rule out B's also being safe. So by taking both "A is safe" and "B is safe" as premises in your practical reasoning, you are not required to cross A, and it remains rational to cross either bridge.
For the delay theory to accommodate this case, the key is to identify the appropriate decision problem. In this example, the decision problem includes three options: to cross either A or B immediately, to go home immediately, or to delay. According Hawthorne and Stanley's view, the first option corresponds to believing the proposition q that "both A and B are safe", while the second option corresponds to disbelieving q. Thus, it is natural to treat delay as suspending judgment on q.
Once the decision problem is framed in an appropriate way, the delay theory can model the situation in which an agent treats more than one believed proposition as a premise for practical reasoning. This allows the delay theory to capture exactly the kind of case Hawthorne and Stanley describe.
To see how the delay theory fits with Weatherson's account, let A be acting immediately and B be delaying. Then Weatherson's view can be stated as follows: one believes that p iff one prefers acting immediately to delaying when q is true, and continues to prefer acting immediately to delaying when $p \vee q$ is true. According to the delay theory, when one rationally suspends judgment on $p \vee q$, one is rational to delay. In this case, since the agent is not delay (assume that the agent is rational), she is not suspending judgment either. Thus, the delay theory also fits with his view.

In particular, when an agent delays, she is waiting to learn more about p so that she can update her doxastic state and decide accordingly.

This feature of delay also directs our attention to an important part of doxastic practice — namely, inquiry — and helps make sense of it. As argued in 4.2, the delay framework naturally captures the connection between suspension and the *action* of inquiry. By contrast, pragmatic encroachment only connects belief to *practical action* rather than epistemic inquiry.

It is then natural to ask what this theory implies about the relationship between belief and delay. Under the tripartite framework of doxastic states, this question is important for fully establishing the theory. Although the topic is worth exploring in depth, it lies beyond the scope of this paper. Owing to space constraints, I develop the theory about belief and delay in a separate companion paper.

5 Conclusion

In conclusion, the delay theory of suspension offers four main advantages over existing credal approaches.

First, it provides unified solutions to the three central challenges facing credal theories: the problems of arbitrary and vague thresholds, the conjunction–disjunction problem, and the tension with the Principal Principle. As del Rio (2024) argues, no precise credal theory can resolve all three, and my analysis suggests how a decision-theoretic framework can meet this challenge.[22] Because suspension is rational exactly when delay is, the thresholds fall where the expected utility of delaying meets that of acting immediately. The framework thereby recovers the view on which belief requires credence 1 (Clarke 2013, Greco 2015) as the limiting case in which delay is costless.

Second, the theory accommodates the inquiry-based understanding of suspension by grounding suspension in the rationality of delay. Third, it aligns naturally with pragmatic encroachment theories by locating belief, disbelief, and suspension within a unified decision-theoretic framework based on delay. Finally, the theory illuminates the epistemic value of suspension and the diachronic epistemic value of belief by assigning epistemic utility to delaying and to maintaining a belief over time, as discussed in Section 4.2.

Given that the theory is built on a tripartite framework of doxastic attitudes, it also invites further questions — not only about suspension, but about familiar problems for credal theories of belief, such as the Lottery and Preface paradoxes. These issues merit separate treatment, and I address them in a companion paper on belief and delay, where I argue that grounding belief in delay can shed new light on these longstanding puzzles.

In sum, the delay theory of suspension offers a principled and flexible credal theory of suspension that preserves the strengths of credal theories while avoiding their common problems. It clarifies the nature of suspension, reveals its epistemic significance, and contributes to broader work on epistemic utility theory and pragmatic encroachment.

[22] Del Rio also raises difficulties for imprecise credal theories, but since this paper develops a precise framework, those discussions lie outside its scope.

References

[1] Atkins, P. (2017). A Russellian account of suspended judgment. *Synthese*, 194(8), 3021–3046.

[2] Basu, R. and Schroeder, M. (2019). Doxastic Wronging. In Kim, B. and McGrath, M. *Pragmatic Encroachment in Epistemology*. New York: Routledge. pp.181-205.

[3] Buchak, L. (2010). Instrumental rationality, epistemic rationality, and evidence-gathering. *Philosophical Perspectives*, 24, 85–120.

[4] Christensen, D. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford, GB: Oxford University Press.

[5] Clarke, R. (2013). Belief Is Credence One (in Context). *Philosophers' Imprint*, 13(11), 1–18.

[6] Crawford, S. (2004). A solution for Russellians to a puzzle about belief. *Analysis*, 64(3), 223–229.

[7] del Rio, A. (2024). Suspending belief in credal accounts. *Noûs*, 58(1), 3-25.

[8] Dorst, K., (2019). Lockeans Maximize Expected Accuracy. *Mind*, 128(509), 175–211.

[9] Dorst, K., and Mandelkern, M. (2022). Good guesses. *Philosophy and Phenomenological Research*, 105(3), 581-618.

[10] Easwaran, K. (2016). Dr Truthlove, Or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Noûs*, 50(4), 816–853.

[11] Fantl, J. and M. McGrath (2002). Evidence, Pragmatics, and Justification. *The Philosophical Review*, 111(1), 67–94.

[12] Fantl, J. and M. McGrath (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

[13] Foley, R. (1992). The Epistemology of Belief and the Epistemology of Degrees of Belief, *American Philosophical Quarterly*, 29(2), 111–121

[14] Foley, R. (1993). *Working Without a Net*. Oxford: Oxford University Press.

[15] Friedman, J. (2013a). Suspended judgment. *Philosophical Studies*, 162(2), 165–181.

[16] Friedman, J. (2013b). Question-directed attitudes. *Philosophical Perspectives*, 27(1), 145–174.

[17] Friedman, J. (2017). Why suspend judging? *Noûs*, 51(2), 302–326.

[18] Friedman, J. (2019). Inquiry and belief. *Noûs*, 53, 296-315.

[19] Good, I. J. (1967). On the principle of total evidence. *The British Journal for the Philosophy of Science*, 17(4), 319–321.

[20] Greco, D. (2015). How I learned to stop worrying and love probability 1. *Philosophical Perspectives*, 29(1), 179–201.

[21] Hájek, A. (1998). Agnosticism meets bayesianism. *Analysis*, 58, 199–206.

[22] Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

[23] Hawthorne, J., and Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571-590.

[24] Hawthorne, J. (2009). The Lockean thesis and the logic of belief. *In Degrees of belief* (pp. 49-74). Dordrecht: Springer Netherlands.

[25] Hawthorne, J., Rothschild, D., and Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5), 1393–1404.

[26] Hempel, C. (1962). Deductive-Nomological vs Statistical Explanation. In H. Feigl and G. Maxwell (Eds.), *Minnesota Studies in the Philosophy of Science, Volume 3* (pp.98–169). Minneapolis: University of Minnesota Press.

[27] Horowitz, S. (2019). Accuracy and educated guesses. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology, Volume 6* (pp. 85–113). Oxford University Press.

[28] Kaplan, D. (2013). De re belief. *The American Philosophical Association Centennial Series*, 9, 25–37.

[29] Lehrer, K. (1970). Justification, explanation, and induction. In M. Swain (Ed.), *Induction, acceptance, and rational belief* (pp. 100–133). D. Reidel.

[30] Leitgeb, H. (2014). The stability theory of belief. *The Philosophical Review*, 123(2), 131–171.

[31] Leitgeb, H. (2017). *The stability of belief: How rational belief coheres with probability*. Oxford University Press.

[32] Levi, I. (1967). *Gambling with truth: An essay on induction and the aims of science*. Alfred A. Knopf.

[33] Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy*, 71(13), 391–418.

[34] Levi, I. (1980). *The enterprise of knowledge: An essay on knowledge, credal probability, and chance.* MIT Press.

[35] Levi, I. (1986). *Hard choices: Decision making under unresolved conflict.* Cambridge University Press.

[36] Masny, M. (2020). Friedman on suspended judgment. *Synthese*, 197(11), 5009-5026.

[37] Moss, S. (2018). IX — Moral encroachment. In *Proceedings of the aristotelian society*, 118(2):177-205. Oxford University Press.

[38] Pettigrew, R. (2024). Epistemic Utility Arguments for Epistemic Norms, *The Stanford Encyclopedia of Philosophy (Summer 2024 Edition)*, Edward N. Zalta and Uri Nodelman (eds.), URL = ¡https://plato.stanford.edu/archives/sum2024/entries/epistemic-utility/¿.

[39] Ross, J. and Schroeder, M. (2014). Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2):259-288.

[40] Salmon, N. (1986). *Frege's puzzle.* Cambridge: MIT Press.

[41] Smith, M. (2010). What Else Justification Could Be. *Noûs*, 44(1), 10–31.

[42] Skyrms, B. (2010). *Signals: Evolution, learning, and information.* Oxford University Press.

[43] Smith, M. (2016). *Between probability and certainty: What justifies belief.* Oxford University Press.

[44] Stalnaker, R. (1974). Pragmatic presuppositions. In Munitz, M. K. and Unger, P., (Eds), *Semantics and Philosophy* (pp. 197–213). New York: New York University Press.

[45] Stalnaker, R. (1984). *Inquiry.* Cambridge, MA: The MIT Press.

[46] Stanley, J. (2005). *Knowledge and Practical Interests.* New York: Oxford University Press.

[47] Sturgeon, S. (2010). Confidence and coarse-grained attitudes. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology, Volume 3* (pp.126–149). Oxford: Oxford University Press.

[48] Sturgeon, S. (2020). *The Rational Mind.* Oxford: Oxford University Press.

[49] Titelbaum, M. G. (2022). *Fundamentals of Bayesian epistemology 1: Introducing credences.* Oxford University Press.

[50] Weatherson, B. (2005). Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives*, 19(1):417–43.

[51] Weatherson, B. (2018). Interests, evidence and games. *Episteme*, 15(3), 329-344.