

Total Evidence and Learning with Imprecise Probabilities

Ruobin Gong

Department of Statistics, Rutgers University, USA

RUOBIN.GONG@RUTGERS.EDU

Joseph B. Kadane

Mark J. Schervish

Department of Statistics, Carnegie Mellon University, USA

KADANE@ANDREW.CMU.EDU

MARK@CMU.EDU

Teddy Seidenfeld

Department of Philosophy and Department of Statistics, Carnegie Mellon University, USA

TEDDY@STAT.CMU.EDU

Abstract

In dynamic learning, a rational agent must revise their credence about a question of interest in accordance with the total evidence available between the earlier and later times. We discuss situations in which an observable event F that is sufficient for the total evidence can be identified, yet its probabilistic modeling cannot be performed in a precise manner. The agent may employ imprecise (IP) models of reasoning to account for the identified sufficient event, and perform change of credence or sequential decisions accordingly. Our proposal is illustrated with three case studies: the classic Monty Hall problem, statistical inference with non-ignorable missing data, and the use of forward induction in a two-person sequential game.

Keywords: Corpus of knowledge, forward induction, IP decision rule, non-ignorable missing data, sufficiency

we mean the information is acquired by the agent, either through their own perception or the aid of measurement instruments. Observational knowledge only constitutes a portion of all knowledge available to the agent. Specifically, the agent’s updated credence at the latter time t_2 should reflect not only the observational knowledge A_{t_2} , but also the epistemic fact that they have learned A_{t_2} .

Denote by $K_{t,M}(F)$ the event that the agent learns at time t by method M that event F obtains. The agent’s *corpus of knowledge* at time t consists of the joint event

$$F \ \& \ K_{t,M}(F). \tag{1}$$

Then, the agent’s credence function at t_2 should be

$$Cr_{t_2}(\cdot) = Cred(\cdot \mid A_{t_1} \ \& \ F \ \& \ K_{t_2,M}(F)),$$

which may or may not agree with the assertion

$$Cr_{t_2}(\cdot) \stackrel{?}{=} Cred(\cdot \mid A_{t_1} \ \& \ F). \tag{2}$$

Indeed, the question-marked equality in (2) will not hold, if the observable event F and its attainment $K_{t_2,M}(F)$ become *epistemologically entangled*. That is, the meaning of the observational report F depends on the context of its attainment, $K_{t_2,M}(F)$, in a non-trivial fashion. In order to update their credence from t_1 to t_2 under epistemological entanglement, the agent will have to specify at the outset a rational credence function in relation to their corpus of knowledge (1), i.e.

$$Cred(F \ \& \ K_{t,M}(F)).$$

This can be a daunting requirement for two reasons. First, the rational credence function $Cred(\cdot)$ needs to be well-defined for all F , t and M . These aspects together span an enormous state space, on which probabilistic specification can be difficult, if not impossible. Second, for a general observable event F , the epistemological information of its attainment $K_{t,M}(F)$ is typically unobservable. Such is true even when granted that the agent satisfies the *KK-thesis*, i.e. whenever they know F , they know that they know it.

1. Introduction

Let $Cr_t(\cdot)$ be an unconditional probability function, called a *rational Credence function* that depicts some idealized agent’s uncertainty at time t . Carnap’s *Principle of Total Evidence* [1] requires that

$$Cr_t(\cdot) = Cred(\cdot \mid A_t),$$

where $Cred(\cdot \mid \cdot)$ is a conditional probability function and A_t is all the observational knowledge that the agent knows at time t . This implies the Bayesian rule of temporal updating, that if between an earlier time t_1 and a later time t_2 the agent’s total knowledge changes by the observational report F , then

$$Cr_{t_2}(\cdot) = Cr_{t_1}(\cdot \mid F).$$

This Carnapian account of idealized Bayesian learning may be limiting. Sometimes, it is difficult to see how the agent’s total evidence at the later time may be represented by a proposition, A_{t_2} , that reports all the *observational* knowledge accumulated prior to t_2 . By observational knowledge,

To circumvent the epistemological entanglement and maintain the feasibility of uncertainty reasoning using probabilities, we argue that the agent may cleverly identify an observable event that nevertheless meets the *Total Evidence* condition, i.e. some special F such that (2) holds with equality. This requires a concept of *sufficiency* of an observable with respect to a corpus of knowledge, put forward by Definition 1.

Definition 1 *An event F observable between times t_1 and t_2 is said to be sufficient for $(F, K_{t_2}(F))$ with respect to a question $\{E, E^c\}$ asked at t_2 , provided that*

$$\text{Cred}(K_{t_2}(F) \mid A_{t_1} \ \& \ F \ \& \ E) = \text{Cred}(K_{t_2}(F) \mid A_{t_1} \ \& \ F). \quad (3)$$

Following this definition, Lemma 2 ensures that when an observable event F is sufficient for the total evidence gained between times t_1 and t_2 , then Carnap's rule of conditionalization may be satisfied in the temporal updating of the agent's credence.

Lemma 2 *If between times t_1 and t_2 the total evidence that the agent gains is the conjunction $(F, K_{t_2}(F))$, then*

$$\text{Cr}_{t_2}(E) = \text{Cr}_{t_1}(E \mid F)$$

if and only if F is sufficient for $(F, K_{t_2}(F))$ with respect to the question $\{E, E^c\}$.

We discuss situations in which the agent is capable of identifying an observable event F that is sufficient for the total evidence, but cannot perform its probabilistic modeling in a precise manner. The identified event F offers more information than a mere observational report the agent can obtain between the earlier and later times. Indeed by sufficiency, F is meant to encode not only the observational report, but also the means through which the agent obtains the report. Therefore, the agent may not have a non-ambiguous probability model to account for F . We utilize imprecise probabilities to analyze an agent's change of credence as a dynamic learning process. In what follows, we illustrate our proposal using three case studies, in the contexts of the classic Monty Hall problem (Section 2), statistical inference with non-ignorable missing data (Section 3), and the use of forward induction in a two-person sequential game (Section 4). Section 5 concludes with a discussion on the operational necessity of our proposal.

2. Monty Hall Problem

In the Monty Hall problem [22, 21, 24], a valuable prize is hidden at random behind one of three enumerated doors: A, B, or C. The other two doors hide no prize. The Contestant makes a first move by designating one of the three doors. The game's moderator Monty Hall then opens one of the other two doors to reveal an empty door. Last, the

Contestant decides whether she would like to stay with the designated door as her final choice, or switching to the third and remaining closed door. She wins if her final choice door hides the prize. Without loss of generality, suppose that the Contestant designated door A as her initial door at t_1 . The question is, what is her credence at t_2 about door A being the prize door, after Monty reveals an empty door to her?

The Contestant knows that the prize was placed uniformly randomly behind one of the three doors, and no further information whatsoever was supplied to her at the first stage of the game. Letting E denote the prize door, we have that the Contestant's credence about E at t_1 is uniform:

$$\text{Cr}_{t_1}(E) = \frac{1}{3}$$

and

$$\text{Cr}_{t_1}(E \mid \text{designate A}) = \frac{1}{3}$$

for all $E \in \{A, B, C\}$. Furthermore, letting D denote the door to be revealed to the Contestant as empty, we have that

$$\text{Cr}_{t_1}(E = A \mid \text{designate A}, D) = \frac{1/3}{1/3 + 1/3} = \frac{1}{2} \quad (4)$$

for both $D \in \{B, C\}$. That is, the Contestant's conditional credence for door A to be the prize door becomes $1/2$, upon knowing either door B or door C is empty.

However, one would be mistaken to think that the Contestant's credence about whether A is the prize door at time t_2 , $\text{Cr}_{t_2}(E = A)$, is represented by the quantity (4). The total evidence available to the Contestant at t_2 is not just that door D is empty, but also the fact that Monty Hall revealed door D to be empty, with $D \in \{B, C\}$. Furthermore, it is through and only through Monty's reveal that the Contestant learns door D to be empty. Therefore, the observable event F that is sufficient (in the sense of Lemma 2) for the Contestant's total evidence is

$$\text{MHRReveals}(D),$$

which in the eyes of the Contestant satisfies

$$\begin{aligned} \text{Cr}_{t_1}(K_{t_2}(D) \text{ iff MHRReveals}(D) \\ \text{iff } K_{t_2}(\text{MHRReveals}(D))) = 1. \end{aligned}$$

Having identified the observable event sufficient for her total evidence, the Contestant's credence about the prize door at time t_2 retains an element of imprecision. In the case the designated door, A, were indeed the prized door, Monty would have the liberty to choose between either door B or door C to reveal to the Contestant, as either door would be empty. As the Contestant has no information about Monty's inclination to reveal either door when he has that choice, her conditional credence function for the joint event $(E, \text{MHRReveals}(D))$ given that she designated door A

is represented by an imprecise probability. Table 1 specifies this imprecise probability, in which E represents the true door behind which the prize stands, D the door that Monty reveals to be empty, and $x \in [0, 1]$ Monty's inclination to reveal door B over door C when he has the liberty to do both. This implies that

$$Cr_{t_1}(E = A \mid \text{designate A, MHRReveals}(D = B)) \quad (5)$$

$$= \frac{x/3}{x/3 + 1/3} \in [0, 1/2], \quad \text{and}$$

$$Cr_{t_1}(E = A \mid \text{designate A, MHRReveals}(D = C)) \quad (6)$$

$$= \frac{(1-x)/3}{(1-x)/3 + 1/3} \in [0, 1/2].$$

Taking into account the total evidence available at t_2 , the Contestant's credence $Cr_{t_2}(E = A)$ is equal to either (5) in case Monty revealed door B to her, or (6) in case Monty revealed door C to her. Therefore, her latter credence $Cr_{t_2}(E)$ is represented by the set of probabilities

$$\mathcal{P} = \{P : P(A) \in [0, 1/2]\},$$

regardless of which door Monty Hall reveals to her.

It is worth noting that in this analysis, since we assume that the Contestant has no information whatsoever about Monty's inclination x , her latter credence exhibits *dilation* [20] when compared to her former credence $Cr_{t_1}(E)$. For the same E , the range of values that $Cr_{t_2}(E)$ may take strictly contains that of $Cr_{t_1}(E)$, regardless of which event in the partition of the total evidence space realizes between t_1 and t_2 , that is, regardless of which (not- E) door Monty Hall reveals to the Contestant. [6] examines dilation in this example in further detail.

Table 1: $Cr_{t_1}(E, \text{MHRReveals}(D) \mid \text{designate A})$, where $E \in \{A, B, C\}$ is the true door to the prize, and $D \in \{A, B, C\}$ the door that Monty Hall Reveals to the Contestant, given that the Contestant designated door A as her initial choice.

Prize door E	Monty Hall Reveals D		
	A	B	C
A	0	$x/3$	$(1-x)/3$
B	0	0	$1/3$
C	0	$1/3$	0

3. Non-ignorable Missing Data

Suppose an experiment is designed to address questions about some feature pertaining to the N members of a population, N being potentially infinite. For each member i

of the population, let X_i denote the true state of their feature. At time t_1 , a simple random sample of n members of the population was surveyed. By time t_2 , however, only $n_{obs} < n$ observations responded, whereas $n_{mis} = n - n_{obs}$ values are missing. Letting X_{obs} denote the collection of n_{obs} observed responses, it is widely understood that the conditional credence

$$Cr_{t_1}(\cdot \mid X_{obs})$$

is not necessarily the correct credence that the investigator should endorse at t_2 . It does not take into consideration the total evidence available to the investigator, which should include the fact that a specific fraction of the sampled members did not respond.

The explicit accounting for the nonresponse requires the introduction of an additional binary observable random variable $D = (D_1, \dots, D_n)$. If the surveyed individual i responded then $D_i = 1$, and $D_i = 0$ if they did not respond. The observed and missing observations can respectively be denoted as

$$X_{obs} = \{X_i : i = 1, \dots, n, D_i = 1\},$$

$$X_{mis} = \{X_i : i = 1, \dots, n, D_i = 0\},$$

and accordingly $n_{obs} = \sum_{i=1}^n D_i$ and $n_{mis} = \sum_{i=1}^n (1 - D_i)$. The investigator's total evidence at time t_2 is

$$(X_{obs}, D). \quad (7)$$

The observable event (7) is sufficient for the investigator's credence for θ if and only if

$$Cr_{t_2}(\theta) = Cr_{t_1}(\theta \mid X_{obs}, D). \quad (8)$$

The assertion (8) lies at the foundation of the missing data literature, and is key to avoiding epistemic entanglement using observable evidence. To update their credence for the scientific question of interest despite partially missing observations, the investigator must be able to supply some kind of knowledge about the nonresponse mechanism. This requirement may well be hard to satisfy. A most challenging type of nonresponse mechanism to model is the *non-ignorable* mechanism [14]. Non-ignorability refers to the case when the response probabilities depend nontrivially on the values of the missing data. By definition, then, any observed and partially missing dataset contains only limited (if any) information about the non-ignorable mechanism. This is precisely why modeling non-ignorability is difficult in practice. The investigator often must conduct post-survey coverage studies in order to gain the needed insight.

As a concrete example, suppose $X_i \in \{0, 1\}$ is a binary feature for an individual, and the investigator is interested in studying θ , the population proportion of individuals possessing the positive feature. Further suppose that the positive feature $X_i = 1$ is associated with an adverse health or

social perception, e.g. that a person smokes. Therefore, an individual who possesses this positive feature is less likely to respond to the survey. (For simplicity, we do assume that if an individual responds, then they respond truthfully.) For a real study on non-ignorability in opinion surveys concerning smoking, see [18].

The investigator posits a parametric sampling model

$$X_i | \theta \sim \text{Ber}(\theta),$$

and a nonresponse mechanism such that for some constant $\gamma \in [0, 1)$,

$$\begin{cases} D_i \sim \text{Ber}(\gamma) & \text{if } X_i = 1, \\ D_i = 1 & \text{if } X_i = 0. \end{cases}$$

That is, all surveyed individuals with a negative feature responded, whereas individuals with a positive feature only respond with probability $\gamma < 1$. This nonresponse mechanism is non-ignorable, because the response indicators D are dependent on the values of the missing data X_{mis} .

To proceed with the analysis, we note that the likelihood function for θ is a marginal likelihood function, integrating out the unobserved missing responses X_{mis} . Writing $s_{\text{obs}} = \sum_{i=1}^n X_i D_i$, the sum of observed positive responses, the likelihood function takes the form

$$\begin{aligned} P(X_{\text{obs}}, D | \theta) &= \sum_{X_{\text{mis}} \in \{0,1\}^{n_{\text{mis}}}} P(X_{\text{obs}}, X_{\text{mis}} | \theta) P(D | X_{\text{obs}}, X_{\text{mis}}) \\ &= (\theta \gamma)^{s_{\text{obs}}} (1 - \theta)^{n_{\text{obs}} - s_{\text{obs}}} [\theta(1 - \gamma) + (1 - \theta)]^{n_{\text{mis}}}. \end{aligned}$$

For the purpose of illustration, suppose that the investigator's prior credence function $Cr_{t_1}(\theta)$ is characterized by the Beta(α, β) family of distributions, with density

$$B^{-1}(\alpha, \beta) \theta^{\alpha-1} (1 - \theta)^{\beta-1},$$

where $B(a, b)$ is the Beta function. Writing $\alpha_{\text{obs}} = s_{\text{obs}} + \alpha$ and $\beta_{\text{obs}} = n_{\text{obs}} - s_{\text{obs}} + \beta$, by (8) we have that the investigator's posterior credence function $Cr_{t_2}(\theta)$ has density

$$f_{\gamma}(\theta) = c_{\gamma}^{-1} \theta^{\alpha_{\text{obs}}-1} (1 - \theta)^{\beta_{\text{obs}}-1} [\theta(1 - \gamma) + (1 - \theta)]^{n_{\text{mis}}}, \quad (9)$$

where the normalizing constant

$$c_{\gamma} = B(\alpha_{\text{obs}}, \beta_{\text{obs}}) \mathcal{R}((\alpha_{\text{obs}}, \beta_{\text{obs}}), (1 - \gamma, 1), -n_{\text{mis}}),$$

where $\mathcal{R}(b, Z, -d)$ is Carlson's multiple hypergeometric function [2], which has been previously studied in the Bayesian modeling of censored categorical data [3, 9] to represent the expectation of marginal linear combinations of Dirichlet random variables. The value of $\mathcal{R}((\alpha_{\text{obs}}, \beta_{\text{obs}}), (1 - \gamma, 1), -n_{\text{mis}})$ is equal to $(1 - \gamma)^{n_{\text{mis}}} {}_2F_1(-n_{\text{mis}}, \beta_{\text{obs}}; \alpha_{\text{obs}} + \beta_{\text{obs}}; \gamma/(\gamma - 1))$, where

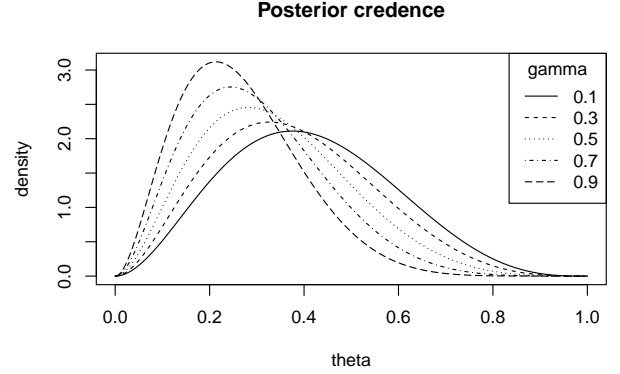


Figure 1: Posterior credence function $Cr_{t_2}(\theta)$ in (9) for different values of γ , the individual response probability with positive feature, for a hypothetical observation with $(n, n_{\text{obs}}, s_{\text{obs}}) = (10, 5, 2)$. Prior credence $Cr_{t_1}(\theta)$ is uniform on $[0, 1]$.

${}_2F_1(u_1, u_2; l_1; z)$ is the generalized hypergeometric function.

The posterior credence function $Cr_{t_2}(\theta)$ depends on the response probability γ for the positive feature. It remains for the investigator to determine what values of γ is reasonable. With γ left unspecified, the posterior credence function $Cr_{t_2}(\theta)$ in (9) induces a set of probability functions

$$\mathcal{P} = \left\{ P : P(A) = \int_A f_{\gamma}(\theta) d\theta, \gamma \in [0, 1) \right\}. \quad (10)$$

Figure 1 depicts $Cr_{t_2}(\theta)$ for different values of γ , for a hypothetical sample with $(n, n_{\text{obs}}, s_{\text{obs}}) = (10, 5, 2)$. Prior credence $Cr_{t_1}(\theta)$ is uniform on $[0, 1]$, corresponding to $\alpha = \beta = 1$. Note that the triple $(n, n_{\text{obs}}, s_{\text{obs}})$ is a reduction of the sufficient observable event in (7), and is *minimally sufficient* for the posterior credence $Cr_{t_2}(\theta)$ in the usual sense of the phrase. As is clear from Figure 1, the posterior credence function $Cr_{t_2}(\theta)$ exhibits large differences depending on the value of γ . If γ is small, it suggests that individuals with $X_i = 1$ are much less likely to respond. Therefore, the fact that half of the surveyed individuals did not respond should be taken as strong indication that there are more people with a positive feature $X_i = 1$ that are unobserved, and the investigator should put higher posterior credence for θ on the larger values. Whereas if γ is large, individuals with $X_i = 1$ are not much less likely to respond, and the posterior credence for θ tend towards the smaller values.

We remark that the IP treatment presented here for the case of non-ignorable missing data has close ties to the literature of partial identification in econometrics; see e.g. Chapter 1 of [17]. Indeed, the investigator's updated credence is partially identified, in the sense that the observed

data (X_{obs}, D) do not provide enough discerning information to pin down $Cr_2(\theta)$ as a unique probability. The identification region of $Cr_2(\theta)$ is precisely the set of probabilities specified by (10). If the investigator would like to avoid partial identification, he or she may adopt a “full Bayesian” approach by further imposing a precise prior credence function on γ , the probability of missing the observation given a positive feature. However, since the observed data do not provide identifying information about γ , the investigator’s future credences about the primary question of interest, θ , may be sensitive to the prior specification for γ . Careful deliberation should be practiced.

4. Forward Induction with Imprecise Probabilities

In this section, we consider two-person sequential games in which each player is required to specify a *grand plan* for action at each node of the game tree where the player needs to make choices. A game-theoretic criterion that dictates what kind of grand plans should be deemed as acceptable is *subgame perfection*; see e.g. [8]. Under subgame perfection, a grand plan is said to be acceptable if and only if it yields acceptable strategies within *each* sub-game of the larger game. However, since players who adhere to subgame perfection must treat each subgame as a separate game irrespective of all other aspects of the larger game, including choices that have been made by their opponents, they can violate total evidence in devising their grand plans for the game.

To take into account the total evidence, the player may endorse instead the criterion of *forward induction* [11], by recognizing and utilizing what they observe from the preceding plays that lead them into the subgame. That is, at the beginning of each subgame, the player works with the total evidence which includes not only their uncertainty model about their opponents, but also the fact that the dynamic of the sequential game has lead both of them to this particular subgame.

In addition to adopting forward induction reasoning in the sequential game, we also assume that the players employ imprecise probability models of uncertainty, represented by a set \mathcal{P} of personal probability functions, rather than by a single such function, P . The literature has demonstrated the value, and indeed the necessity, of using IP models of uncertainty in game theory [e.g. 5, 23, 15].

We now describe the setting of the game. Two players, Sidney and Isaac, are about to play a two-stage extensive form sequential game. In the first stage of the game, Sidney chooses between their playing either the Concert Game, or playing the Lecture Game. In the Concert game, they coordinate on attending either (A) a Bruch violin concerto, played by Itzhak Perlman, or (B) a Dolly Parton concert. In the Lecture game, they coordinate on attending (C) a

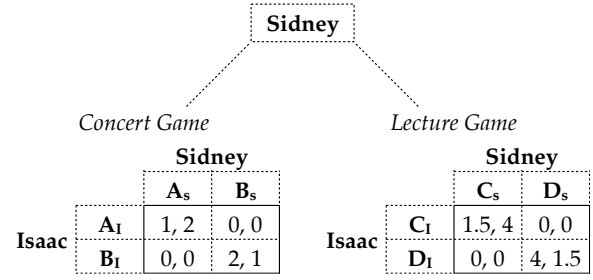


Figure 2: A two-person, two-stage extensive form sequential game, in which one player (Sidney) chooses which subgame for both players to play at the second stage.

lecture by Chomsky, or (D) a lecture by Ellsberg. In these subgames, Sidney is Column player and Isaac is Row player. The goal of both players is to arrive at a precise action plan at the end of the iterative game, while at the same time maximizing their own utility outcome to the extent possible. We pay particular attention to the interpretation of players’ strategies and the utility outcomes from Isaac’s point of view, whose application of forward induction highlights the impact of total evidence in guiding his subgame decision as the Row player.

In either the Concert or Lecture subgames, we allow the two players to adopt an extreme IP model that reflects maximal uncertainty about the other player’s choices. That is in each subgame, each player uses the set \mathcal{P} of all probabilities for what the other player might choose from among all his mixed strategies. For instance, without additional evidence, Isaac is maximally uncertain about which strategy Sidney will use in the Lecture Game. That is,

$$\begin{aligned} \mathcal{P}_{\text{Lecture}}^{\text{Isaac}} \{xA_S \oplus (1-x)B_S\} \\ = \{P : 0 \leq P(xA_S \oplus (1-x)B_S) \leq 1\}, \forall x \in [0, 1]. \end{aligned}$$

The two subgames each have the following three Nash equilibria pairs, two are pure and one is mixed. In the Concert Game:

- $\langle A_I, A_S \rangle$ yields utility outcome (1, 2);
- $\langle B_I, B_S \rangle$ yields utility outcome (2, 1);
- $\langle (1/3)A_I \oplus (2/3)B_I, (2/3)A_S \oplus (1/3)B_S \rangle$ yields utility outcome (2/3, 2/3).

In the Lecture Game:

- $\langle C_I, C_S \rangle$ yields utility outcome (1.5, 4);
- $\langle D_I, D_S \rangle$ yields utility outcome (4, 1.5);
- $\langle (3/11)C_I \oplus (8/11)D_I, (8/11)C_S \oplus (3/11)D_S \rangle$ yields utility outcome (12/11, 12/11).

One complication with the analysis of games represented by IP models is that there exists a variety of applicable decision rules. A different choice of rules may yield different action plans and different consequences [19, 16]. In the current two-stage game, both players aim to arrive at one precise action plan, and the IP decision rule that they employ must be conducive to this goal. Thus in this example, the IP decision rule that both players employ restricts admissibility to those options that maximize minimum expectation with respect to the set \mathcal{P} of probabilities, i.e. options that are Γ -maximin [4], among those options that maximize expected utility for some probability P in the set \mathcal{P} , i.e. options that are E -admissible [12]. This is Levi's lexicographic rule [13] that uses E -admissibility as the primary consideration, and Γ -maximin as the secondary (or *security*) consideration for admissibility. A brief discussion about the choice of IP decision rules appears at the end of this section.

Since each player has a maximally uncertain IP model for which strategy the other player chooses in these games, each of these three Nash pairs also are pairs of E -admissible options, because each of these three strategies maximizes expected utility against the other's matching strategy. More significant, in each game, in the light of the "equalize" mixed strategy Nash equilibrium pair, *each* mixed strategy that Isaac might choose also is E -admissible against that mixed strategy for Sidney. For instance, in the Lecture Game, each mixed strategy that Isaac might play, $x C_I \oplus (1-x) D_I$, is E -admissible against Sidney's equalizer mixed strategy, $(8/11) C_S \oplus (3/11) D_S$.

Next, we turn to considerations of security maximization. In the Concert Game:

- Isaac's mixed strategy, $(2/3) A_I \oplus (1/3) B_I$ secures a minimum expectation of $2/3$, and that is the maximum security possible for Isaac among his (E -admissible) strategies.
- Likewise, by the symmetries of the game, Sidney's mixed strategy $(1/3) A_S \oplus (2/3) B_S$ secures a minimum expectation of $2/3$, and that is the maximum security possible for Sidney relative to all his (E -admissible) strategies.

In the Lecture Game, the security maximizers are

- For Isaac, $(8/11) C_I \oplus (3/11) D_I$ secures $12/11$ utility, and
- For Sidney, $(3/11) C_S \oplus (8/11) D_S$ secures $12/11$ utility.

Here is how we apply forward induction with these IP decision rules in the two-stage game between Isaac and Sidney, where Sidney plays first to choose which subgame they play. We use the following hypothetical "cheap talk" dialogue to make explicit the steps in the IP decision making.

Sidney: Isaac, suppose I choose we to go to the Concert. What will you do?

Isaac mumbles to himself: Well, if I saw that Sid chose the Lecture Game, that would give him an E -admissible option with a security of $12/11$. Hmm...

Isaac: Then, Sid, if you choose the Concert Game (and reject the Lecture Game) you'd be signaling to me that you expect at least $12/11$ in the Concert Game. So, I'd choose to join you to hear Perlman play Bruch, and you'll get 2 units utility while I get only 1.

Sidney: Very good. Let's go to the Lecture!

Isaac mumbles to himself: Well, rejecting Concert means that Sid now expects at least 2 units by going to the Lecture.

Isaac: Then, Sid, I see I'm stuck going to hear Chomsky with you.

Sidney: Yes. But at least you'll enjoy that more than you would the Bruch!

Note, the application of forward induction illustrated in this example conforms to the conjecture that players (e.g. Isaac) can avoid the epistemic entanglement by using observable (even hypothetical) decisions from earlier in the game to fix expectations later in the game, without needing to incorporate an additional epistemic random variable for current knowledge. Indeed, in the first iteration of the game, Isaac chooses Perlman's Bruch (A_I) over his preferred Dolly Parton (B_I), because he knows that Sidney rejected altogether the Lecture subgame (and thereby a security of $12/11$), a piece of observed knowledge that precedes the current Concert subgame. Similarly, in the second iteration of the game Isaac chooses Chompsky (C_I) over his preferred Ellsberg (D_I), because he knows that Sidney rejected the Bruch concerto (and thereby a certain utility outcome of 2) in the Concert subgame, which is again observed knowledge that precedes the current Lecture subgame.

Before concluding this section, we remark on the use of Levi's lexicographic IP decision rule in this example. The primary purpose of the example is to illustrate forward induction in sequential games with ambiguity, as a means for the players to avoid the epistemic entanglement using sufficient observables. Given that both players possess vacuous knowledge about each other's strategies, this lexicographic decision rule (i.e. E -admissibility first, with Γ -maximin as secondary security) allows the players to arrive at a unique strategy. As discussed, E -admissibility alone reduces the admissible options only to the infinite number of *rationalizable* strategies. On the other hand, if both players endorse Γ -maximin without consideration for E -admissibility, it would hinder Isaac's ability to perform forward induction and make use of Sidney's suggestion for the Concert game as evidence to guide his own choice. In

Isaac’s view, Sidney’s strategy needs not be Nash, if he is not bound by E -admissibility.

We do not defend the lexicographic rule as the “correct” rule for this game. Nor do we preclude the possibility that other IP decision rules may offer sensible alternative analyses that deliver a unique strategy at the end, and to help the players avoid the epistemic entanglement. In fact, one may question the merits of the lexicographic rule, on the grounds of information value. It is understood that the lexicographic rule does not necessarily respect the value of cost-free, new information [19]. The mere suggestion by Sidney that they might play the Concert game, despite being a hypothetical one, is enough to steer the game towards the unique outcome that maximizes Sidney’s utility globally, but not Isaac’s. The answers to some questions remain open for further research. For example: (i) What was Isaac’s assessment on the *net value* [10] of his total evidence? and (ii) In general, what should a player do when their total evidence incurs a negative net value?

5. Discussion

We have discussed how imprecise probabilities can help an agent to update their temporal credence with respect to the total evidence, in case a sufficient and observable reduction to it can be found. A question that could have been asked in the first place is whether the sufficiency requirement is necessary. In other words, instead of worrying about finding a sufficient observable event F to serve as a reduction, why don’t we consider IP models for the total evidence pair $(F, K_{t,M}(F))$ as in (1)?

The kind of imprecise probabilities that we employ in this paper may not be able to capture *all* varieties of uncertainty and ignorance that a rational agent may have. An IP model is a collection of probabilities defined on a common state space associated with a common sigma field. IP models are useful when the agent is unable to pinpoint their credence function in relation to their corpus of knowledge, nevertheless seeks to update their credence as new information is learned. However, the agent must be certain about their corpus of knowledge, for it is the basis on which to derive any credence at all, precise or otherwise. Expression of uncertainty that pertains to the act of knowing, such as captured by the phrase “I’m not sure if I know F ”, calls for constructs such as *probabilities of higher types* as advocated by Jack Good [7]. In contrast to IP models, however, higher types of probabilities are fuzzy not only in themselves, but also in the inequalities that can express them. They pose a different challenge in terms of their operationalization, and are therefore out of scope for this paper.

Another open question is whether it is always possible for the agent to find an observable sufficient reduction to their total evidence. We surmise the answer may not be categorically affirmative. In our formulation of the agent’s

corpus of knowledge (1), the event $K_{t,M}(F)$ that signifies the attainment of the observational report F depend not only on the time of observation t , but also the method M through which the observation of F can be made. For certain method M , to ascertain the event $K_{t,M}(F)$, or any sufficient reduction of it, may well be infeasible for the agent. For example, the measurement of certain complex scientific phenomena is viable only in theory, or may be too costly to perform. Nevertheless, if the agent can identify an affordable and practically observable sufficient reduction for which only ambiguous credence is available, they should prefer it to an unattainable one for which precise credence is available. As demonstrated in this paper, the agent may avoid the epistemic entanglement and extract meaningful inference from the former, using the tools of imprecise probabilities.

Acknowledgments

We thank three anonymous reviewers for their helpful comments. The research of R. Gong is supported in part by the National Science Foundation (DMS-1916002).

References

- [1] Rudolf Carnap. The aim of inductive logic. In Ernest Nagel, Patrick Suppes, and Alfred Tarski, editors, *Logic, methodology and philosophy of science*, chapter 5, pages 303–318. Stanford University Press, Stanford, CA, 1962.
- [2] James M Dickey. Multiple hypergeometric functions: Probabilistic interpretations and statistical uses. *Journal of the American Statistical Association*, 78(383): 628–637, 1983.
- [3] James M Dickey, Jhy-Ming Jiang, and Joseph B Kadane. Bayesian methods for censored categorical data. *Journal of the American Statistical Association*, 82(399):773–781, 1987.
- [4] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153, 1989.
- [5] Raphaël Giraud. Objective imprecise probabilistic information, second order beliefs and ambiguity aversion: an axiomatization. In *Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications (ISIPTA’05)*, pages 183–192, 2005.
- [6] Ruobin Gong and Xiao-Li Meng. Judicious judgment meets unsettling updating: dilation, sure loss, and Simpson’s paradox (with discussion). *Statistical Science*, 36(2):169–190, 2021.

- [7] Irving John Good. Symposium on current views of subjective probability: Subjective probability as the measure of a non-measurable set. In *Studies in Logic and the Foundations of Mathematics*, volume 44, pages 319–329. Elsevier, 1966.
- [8] John C Harsanyi and Reinhard Selten. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.
- [9] Thomas J Jiang, Joseph B Kadane, and James M Dickey. Computation of Carlson’s multiple hypergeometric function R for Bayesian applications. *Journal of Computational and Graphical Statistics*, 1(3): 231–251, 1992.
- [10] Joseph B Kadane, Mark Schervish, and Teddy Seidenfeld. Is ignorance bliss? *The Journal of Philosophy*, 105(1):5–36, 2008.
- [11] Elon Kohlberg and Jean-François Mertens. On the strategic stability of equilibria. *Econometrica*, 54(5): 1003–1037, 1986.
- [12] Isaac Levi. On indeterminate probabilities. *The Journal of Philosophy*, 71(13):391–418, 1974.
- [13] Isaac Levi. *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. MIT press, 1980.
- [14] Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*. John Wiley & Sons, 2019.
- [15] Hailin Liu. Common knowledge, ambiguity, and the value of information in games. In *Proceedings of the Ninth International Symposium on Imprecise Probability: Theories and Applications (ISIPTA’15, Pescara)*, pages 167–176, 2015.
- [16] Hailin Liu and Wei Xiong. Dynamic consistency in incomplete information games under ambiguity. *International Journal of Approximate Reasoning*, 76: 63–79, 2016.
- [17] Charles F Manski. *Partial identification of probability distributions*. Springer Science & Business Media, 2003.
- [18] Louis T Mariano and Joseph B Kadane. The effect of intensity of effort to reach survey respondents: A Toronto smoking survey. *Survey Methodology*, 27(2): 131, 2001.
- [19] Teddy Seidenfeld. A contrast between two decision rules for use with (convex) sets of probabilities: Γ -maximin versus E -admissibility. *Synthese*, 140(1/2): 69–88, 2004.
- [20] Teddy Seidenfeld and Larry Wasserman. Dilation for sets of probabilities. *The Annals of Statistics*, 21(3): 1139–1154, 1993.
- [21] Steve Selvin. On the Monty Hall problem (letter to the editor). *The American Statistician*, 29(3):134, 1975.
- [22] Steve Selvin. A problem in probability (letter to the editor). *The American Statistician*, 29(1):67, 1975.
- [23] Matthias CM Troffaes. Decision making under uncertainty using imprecise probabilities. *International journal of approximate reasoning*, 45(1):17–29, 2007.
- [24] Marilyn vos Savant. Game show problem, retrieved 02/14/2021. URL <https://web.archive.org/web/20120429013941/http://marilynvossavant.com/game-show-problem/>.