

Available online at www.sciencedirect.com**ScienceDirect**

Journal of Economic Theory ●● (●●●●) ●●●●●●

JOURNAL OF
**Economic
Theory**www.elsevier.com/locate/jet

Updating confidence in beliefs [☆]

Brian Hill*GREGHEC, HEC Paris & CNRS, France*¹

Received 12 September 2019; final version received 25 November 2020; accepted 29 January 2021

Abstract

This paper develops a belief update rule under ambiguity, motivated by the maxim: in the face of new information, retain those conditional beliefs in which you are more confident, and relinquish only those in which you have less confidence. We provide a preference-based axiomatisation, drawing on the account of confidence in beliefs developed in Hill (2013). The proposed rule constitutes a general framework of which several existing rules for multiple priors (Full Bayesian, Maximum Likelihood) are special cases, but avoids the problems that these rules have with updating on complete ignorance. Moreover, it can handle surprising and null events, such as crises or reasoning in games, recovering traditional approaches, such as conditional probability systems, as special cases.

© 2021 Elsevier Inc. All rights reserved.

JEL classification: D81; D83

Keywords: Belief update; Ambiguity; Multiple priors; Confidence; Complete ignorance; Update on surprising or null events

‘[W]e were seeing things that were 25-standard deviation moves several days in a row’ David Viniar, CFO Goldman Sachs (Financial Times, 13 August 2007).

[☆] I wish to thank Richard Bradley and Itzhak Gilboa, as well as audiences at RUD and D-TEA, the editor and referees for stimulating discussions and valuable feedback on early versions of this paper. The author gratefully acknowledges support from the ANR project DUSUCA (ANR-14-CE29-0003-01).

E-mail address: hill@hec.fr.

¹ 1 rue de la Libération, 78351 Jouy-en-Josas, France.

<https://doi.org/10.1016/j.jet.2021.105209>

0022-0531/© 2021 Elsevier Inc. All rights reserved.

1. Introduction

Reasons for going beyond the Bayesian representation of beliefs by probability measures abound. Whether it be decision makers' observed non-neutrality to ambiguity (Ellsberg, 1961), the purported unjustifiability of the Bayesian requirement of belief precision (Gilboa et al., 2009, 2012; Bradley, 2014) or the difficulty of forming warranted beliefs satisfying the Bayesian tenets in real decisions (Cox, 2012; Gilboa and Marinacci, 2013), many have argued for non-probabilistic representations of belief. But how should such non-Bayesian beliefs be updated?

Two sorts of situations pose particular challenges for the update of non-Bayesian beliefs: complete ignorance and surprising events. A venture capitalist faced with an entirely new drug, based on a never-tested, revolutionary technology, is in a situation of *complete ignorance*. In terms of the popular multiple prior model (Gilboa and Schmeidler, 1989), her beliefs concerning the probability of cure under the drug can be characterised by the set of all (relevant) priors, i.e. the set $[0, 1]$. Yet a central determinant of her investment decision will be how she updates on the basis of early tests giving cure rates on relatively small samples. For an example of *surprising events*, consider an investor wedded to a fine-tuned model of the stock market that has great difficulty explaining, say, the events in a financial crisis—the opening quote, given four days after BNP Paribas suspended redemptions from three funds, is an example from the 2008 financial crisis. Such an investor typically has precise probabilistic beliefs about future asset returns and prices (conditional on their past values), which give extremely low probability to transpired events; the issue is how to update them. Going by the growing body of evidence that ambiguity increases in a financial crisis (e.g. Caballero and Krishnamurthy, 2008; Ilut and Schneider, 2014), a common reaction is to keep an open mind about potential alternative models, and accordingly revert to more open sets of probabilities over future asset prices.

Despite the importance of these situations for motivating non-Bayesian beliefs (e.g. Gilboa et al., 2009; Levi, 1974), existing accounts of update struggle to adequately capture learning in them. On the one hand, the main update rules for multiple priors tend to deal with complete ignorance in an 'extreme' way (Gilboa and Marinacci, 2013), for instance by retaining the completely ignorant $[0, 1]$ priors in the previous example, or jumping directly to a perfectly precise posterior belief (see Section 4.1). Neither seems to reflect how a typical venture capitalist would, or should, react. On the other hand, multiple prior update rules typically coincide with Bayesian conditionalisation when priors are precise, and hence suffer from the same issues concerning updating on surprising events—and *a fortiori* on null events—which plagues the Bayesian account (e.g. De Bondt and Thaler, 1985, 1987). They tell the investor in the opening quote to stick to his original model in the face of the conflicting evidence, rather than keeping an open mind.

The current paper takes up this double challenge, developing and behaviourally characterising a novel account of the update of non-Bayesian beliefs. It is, to our knowledge, the first to cope comfortably with both complete ignorance and surprising events (see Table 1). Moreover, it is general enough to recover 'classic' multiple prior update rules, such as Full Bayesian and Maximum Likelihood update, as well as more recent suggestions, including that used by Epstein and Schneider (2007), as special cases. As such, it provides a unified analysis of them. It also recoups existing approaches to updating on null events, such as Conditional Probability Systems, as a special cases, whilst preserving a strong relationship between *ex ante* and *ex post* preferences, which they lack. In so doing, our account provides new perspectives on these existing approaches. Finally, the account also has solid normative credentials, being built on a reasonable story about how beliefs should be updated in the face of contrary evidence.

Conceptually, our account taps into an intuition as to *why* beliefs may be non-Bayesian: decision makers may be more or less *confident* in different beliefs. This ‘second-order’ aspect is something that the Bayesian model has trouble rendering properly, whilst it can be captured, and related to preferences, in some non-Bayesian models (e.g. Hill, 2019). Formally, we represent confidence in beliefs using *confidence rankings* (Hill, 2013). A confidence ranking is a nested family of sets of probability measures, where different sets are understood as representing the beliefs, or probability judgements,² held by the decision maker at different levels of confidence. The larger a set of probability measures, the fewer probability judgements hold for all measures in it; so *larger* sets in a confidence ranking involve *fewer* beliefs in this sense, and accordingly correspond to *higher* levels of confidence. Structures of this sort have long been employed in econometrics (e.g. Manski and Nagin, 1998; Manski, 2013).

Whilst previous work has connected confidence to preferences (Hill, 2013, 2019), the account developed here recognises that it *also* has a role to play in update. Put succinctly: in updating beliefs, *retain* those conditional beliefs in which you are *more* confident, and relinquish only those in which you have less confidence. Formally, this is reflected in the following update rule (see Section 2.4 for details): given prior confidence ranking Ξ , on learning the event E , the posterior confidence ranking is

$$\Xi_E = \left\{ \{p \in \mathcal{C} : p(E) \geq \rho_E(\mathcal{C})\}_E : \mathcal{C} \in \Xi, \{p \in \mathcal{C} : p(E) \geq \rho_E(\mathcal{C})\} \neq \emptyset \right\} \quad (1)$$

where $\rho_E : \Xi \rightarrow [0, 1]$ is a decreasing function and, for every set of probability measures \mathcal{C} and event E , \mathcal{C}_E is the well-known Full Bayesian update defined as follows:

$$\mathcal{C}_E = \{p(\bullet/E) : p \in \mathcal{C}, p(E) > 0\} \quad (2)$$

The *probability-threshold* function ρ_E assigns a probability value to every set in the confidence ranking, and hence implicitly to every confidence level. In so doing, it effectively specifies a set of probability measures for each confidence level, namely those which assign ex ante probability to E greater than the ρ_E -value for that level. These can be thought of as representing the conclusions the decision maker is warranted to deduce from the observation of E with that much confidence: any probability measure giving a value to E that is less than this threshold ‘gets it too wrong’ to be considered plausible at that confidence level. Since ρ_E is decreasing, this set is larger for higher confidence levels: the conclusions that can be drawn from the data with high confidence are weaker than those that can be drawn with lower confidence. For instance, the observation of drug trials on 100 patients, out of which 75 were cured, warrants high confidence that the probability of cure for a randomly selected new patient is 0.25 or higher, but more limited confidence that this probability is 0.60 or higher. Probability thresholds are reminiscent of significance levels in hypothesis testing, and indeed there is a sense in which the proposed update rule retains the spirit of classical statistical reasoning (see Section 5).

Update rule (1) is visually illustrated on Fig. 1. For every confidence level, the prior beliefs held at that level are represented by the appropriate set of probability measures in Ξ , whereas the conclusions that can be drawn from the data with that level of confidence are summarized by the set of probability measures singled out as ‘reasonable’ by ρ_E . If these are compatible—if the two sets of probability measures overlap—the update rule (1) retains all of these as posterior

² By probability judgement, we mean a statement concerning probabilities, such as ‘the probability of event E is greater than p ’.

beliefs at that confidence level—it takes the intersection. This corresponds to the maxim that conditional beliefs held or conclusions drawn with high confidence are, as far as possible, retained. By contrast, at lower confidence levels, where the (more precise) initial beliefs may contradict the (stronger) conclusions drawn from the data with that much confidence, neither are retained. In cases of conflict with observation, it is the beliefs held with low confidence, if any, that are withdrawn. Finally, the update rule conditions the structure obtained on the learnt event E .

The update rule can be illustrated on the ‘surprising events’ investor example above. Suppose that the investor was using a model specification—say, a set of dynamic stochastic equations— f and a vector of relevant parameter values α to determine the stochastic process relating future vectors of asset returns and prices z_{t+1}, \dots, z_τ to previous ones z_0, \dots, z_t .³ Letting, for any model specification f' and parameter values α' , $p_{\alpha'}^{f'}$ be the conditional probability measure over future values of z in solution,⁴ the investor’s initial beliefs are represented by p_α^f . However, this gives no information about how confident he is in these beliefs, nor about his relative confidence in the model specification f as compared to the parameter values α . Suppose, for the purposes of illustration, that he is more confident in the parameter values α than in the model specification f . His confidence in beliefs can thus be modelled by a three-level confidence ranking, $\Xi = \left\{ \left\{ p_\alpha^f \right\}, \left\{ p_{\alpha'}^{f'} : f' \in \mathcal{G} \right\}, \left\{ p_{\alpha'}^{f'} : f' \in \mathcal{G}, \alpha' \in \mathcal{B} \right\} \right\}$, where \mathcal{G} is a set of models containing f and \mathcal{B} is a set of possible parameter settings containing α .⁵ According to the bottom element of this confidence ranking, asset returns and prices are determined by the model f with parameters α : this captures the stated ex ante beliefs, independently of the confidence with which they are held. At the next level up, all probability measures in the set correspond to the same parameter values α , but different model specifications. This captures a judgement that he is more confident in the former than the latter. Finally, neither the belief about the model nor that concerning the parameters are retained at the highest confidence level. This confidence ranking is drawn in black in Fig. 1.

On observing an economic event E (e.g. z_{t+1} or a set of possible values for $z_{t+1}, \dots, z_{t'}$, given the sequence z_0, \dots, z_t), the investor can assign to each confidence level a probability threshold, determining which measures can be ruled out as ‘having got the prediction about E too wrong’ with that much confidence. For instance, he could apply thresholds of 0.05, 0.01 and 0.001 for the low, medium and high confidence levels respectively.⁶ The sets of probability measures giving a probability to the observation higher than the threshold are shown in red in Fig. 1.

The confidence ranking resulting from update rule (1) is indicated by the blue shaded sets in Fig. 1.⁷ In this example, at the top two confidence levels, the intersection of the sets is non-

³ In the interests of simplicity, we assume that α covers parameters that are relevant across different model specifications.

⁴ I.e. $p_{\alpha'}^{f'}(z_{t+1}, \dots, z_\tau | z_0, \dots, z_t)$ is the probability of values z_{t+1}, \dots, z_τ given z_0, \dots, z_t , under model f' and parameters α' .

⁵ For instance, for $\alpha \in \mathbb{R}^n$, \mathcal{B} could be $\{\alpha' : d(\alpha, \alpha') \leq \zeta\}$ for some appropriate quasi-metric d and threshold ζ , and \mathcal{G} could be a set of model specifications by dynamic stochastic equations incorporating a variety of stochastic assumptions. Whilst we consider a three-level confidence ranking for simplicity, varying amounts of confidence in ranges of parameter values or model specification assumptions can be represented using richer confidence rankings.

⁶ I.e. $\rho_E(C_{low}) = 0.05$, $\rho_E(C_{medium}) = 0.01$, $\rho_E(C_{high}) = 0.001$.

⁷ More specifically, this is the confidence ranking $\left\{ \left\{ p_{\alpha'}^{f'} : f' \in \mathcal{G}, p_{\alpha'}^{f'}(E | z_0, \dots, z_t) \geq \rho_E(C_{medium}) \right\} \left\{ p_{\alpha'}^{f'} : f' \in \mathcal{G}, \alpha' \in \mathcal{B}, p_{\alpha'}^{f'}(E | z_0, \dots, z_t) \geq \rho_E(C_{high}) \right\} \right\}$.




Fig. 1. Confidence Update. Black sets: the sets in the prior confidence ranking $\Xi = \{C_{low}, C_{medium}, C_{high}\}$; Red sets: the sets of probability measures meeting the probability threshold $\{p : p(E) \geq \rho_E(C_i)\}$ for $i = \{low, medium, high\}$. (To visualise the colours in the figure, the reader is referred to the web version of this article.)

empty, and these yield the posterior beliefs. It follows that all the prior conditional beliefs are retained—for instance, at the medium confidence level, the belief in parameter values α is retained. Perhaps new beliefs are added—in the example, the posterior beliefs at the medium confidence level are more precise than they were ex ante, corresponding to the investor retaining some alternative model specifications as potentially more reasonable. The probability measure at the bottom confidence level in the prior confidence ranking assigns a probability to E that is below the corresponding threshold, so the beliefs specific to that level are dropped on learning. In the face of what, under his prior beliefs about the model and parameters, is a 25-standard deviation event of the sort in the opening citation, the investor retracts those beliefs in which he is least confident: in this case, those concerning the model specification. The only beliefs held at the low confidence level ex post are those inherited from higher confidence levels.

Note that the posterior beliefs held at the lowest confidence level will typically not correspond to a single probability measure, or a single model specification-parameter value pair, but will rather remain open between several competitive models. In belief representation terms, the posterior beliefs are not precise, though the prior ones are (at the low confidence level). So this update rule captures the reaction of ‘keeping an open mind’ after surprising events. Note also that it is the investor’s *ex ante* confidence (represented by the confidence ranking) and its interaction with his confidence in the conclusions drawn from the observation (represented by the probability-threshold function) that determine which beliefs are retained. Since the investor in this example has more confidence in the parameter values, this belief is retained—and the belief about the model specification dropped—on update. This is perfectly coherent with the previously stated maxim exhorting him to retain beliefs in which he is more confident, at the price of those held with less confidence.⁸ This example thus illustrates how the proposed approach can cope with surprising events.

In this paper, we provide a behavioural characterisation of a generalisation of (1) that we call (*general*) *confidence update*. All the parameters, and in particular that playing the role of ρ_E , are revealed from preferences. Special cases, including a version of rule (1), are also axiomatised.

⁸ By contrast, under a prior confidence ranking reflecting an investor who is more confident in the model specification than the parameter values, the model would be retained and beliefs about the parameter values withdrawn on update. And under different probability thresholds, which are more permissive at lower confidence levels, both the model and the parameters would be retained, as under the Bayesian approach.

Table 1
Update rules and the double challenge: a summary.
–: cannot be applied; \times : problematic; \checkmark : adequate; \sim : partially adequate (see footnote).
In brackets: Sections where the rules are discussed.

	Complete ignorance	Surprising events
Bayesian update	–	\times (4.2)
Full Bayesian	\times (4.1)	\times (4.2)
Maximum Likelihood	\times (4.1)	\times (4.2)
Epstein and Schneider (2007)	\checkmark (4.1)	\times (4.2)
Conditional Probability Systems	–	\sim^9 (4.2.3, 5)
Confidence update	\checkmark (4.1)	\checkmark (4.2)

We show that this update rule can deal not only with surprising events but also with updating on complete ignorance. This is summarised in Table 1,⁹ which also lists a range of existing update rules that can be recovered as special cases of confidence update.

The paper is organised as follows. Section 2 sets out the framework, the confidence model and update rules. Section 3 contains the main results of the paper, characterising general and specific versions of confidence update, and considering its comparative statics. Section 4 brings out the contributions of the proposed approach with respect to the issues of update from complete ignorance and on surprising or null events, including null events in game-theoretical reasoning. Section 5 discusses extensions and *inter alia* the relationship to Bayesian and classical statistical reasoning. Proofs and other material are to be found in the Appendix.

2. Preliminaries

2.1. Setup

Let \mathcal{S} be a non-empty set of states, with a σ -algebra Σ of subsets of \mathcal{S} , called *events*. $\Delta(\Sigma)$ is the set of finitely-additive probability measures over (\mathcal{S}, Σ) endowed with the weak* topology. For every subset $\mathcal{C} \subseteq \Delta(\Sigma)$, $\bar{\mathcal{C}}$ denotes the closure of \mathcal{C} . Let \mathcal{X} , the set of *consequences*, be a convex subset of a vector space; for instance it could be the set of lotteries over a set of prizes, as in the Anscombe and Aumann (1963) setting. \mathcal{A} is the set of (*simple*) *acts*: finite-valued Σ -measurable functions from states to consequences. \mathcal{A}^c is the set of constant acts (acts taking a constant value). Mixtures of acts are defined pointwise as standard: for any $f, g \in \mathcal{A}$ and $\alpha \in [0, 1]$, the α -mixture of f and g , which we denote by $f_\alpha g$, is defined by $f_\alpha g(s) = \alpha f(s) + (1 - \alpha)g(s)$ for all $s \in \mathcal{S}$. For every $f, g \in \mathcal{A}$ and $E \in \Sigma$, $f_E g \in \mathcal{A}$ is defined by $f_E g(s) = f(s)$ if $s \in E$, $f_E g(s) = g(s)$ otherwise.

We use \succeq (perhaps with subscripts) to denote a preference relation on \mathcal{A} . The symmetric and asymmetric parts of \succeq , \sim and \succ , are defined as standard. We say that \succeq is *degenerate* if $f \sim g$ for all $f, g \in \mathcal{A}$. A functional $V : \mathcal{A} \rightarrow \mathbb{R}$ is said to represent preferences \succeq if $f \succeq g$ if and only if $V(f) \geq V(g)$.

Henceforth, \succeq (with no subscript) will denote the decision maker's ex ante preferences. Ex post preferences will be denoted with subscripts, depending on the information received; there is a class of preferences $\{\succeq_E\}_{E \in \Sigma}$. For each event E , \succeq_E is the decision maker's preference after

⁹ Note on Conditional Probability Systems entry: Only works for null events; no continuous treatment of null and surprising events (Section 4.2.3). Ex ante preferences place very few restrictions on ex post ones (Section 5).

having learnt (only) that E obtains. Finally, an event $E \in \Sigma$ will be said to be null if it is null with respect to \succeq : for any $f, g, h, h' \in \mathcal{A}$, $f_{E^c}h \succeq g_{E^c}h'$ if and only if $f \succeq g$.

2.2. Confidence model

We adopt the confidence framework set out and developed in Hill (2013, 2016, 2019). Beliefs are represented by a *confidence ranking* on \mathcal{S} : a nested family of non-empty subsets of $\Delta(\Sigma)$. (I.e. a confidence ranking Ξ is a subset of $2^{\Delta(\Sigma)} \setminus \emptyset$ such that, for all $\mathcal{C}, \mathcal{C}' \in \Xi$, $\mathcal{C} \subseteq \mathcal{C}'$ or $\mathcal{C}' \subseteq \mathcal{C}$.) Different sets in the confidence ranking represent beliefs held with different levels of confidence. Note that a single set of probability measures à la Gilboa and Schmeidler (1989) is a degenerate special case of a confidence ranking; it can be interpreted as the case where the same beliefs are held at all levels of confidence.

A confidence ranking Ξ is said to be *closed* (resp. *convex*) if each set in the family is. We let $\min \Xi = \bigcap_{\mathcal{C}' \in \Xi} \mathcal{C}'$ and $\max \Xi = \bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'$; these can be loosely thought of as the bottom and top elements of Ξ . For a confidence ranking Ξ , its min-closure, $\Xi^{\min} = \Xi \cup \{\min \Xi\}$.¹⁰ Ξ is *min-closed* if $\Xi = \Xi^{\min}$. Throughout the axiomatic treatment (Section 3), we shall only be concerned with closed, convex and min-closed confidence rankings.¹¹

As discussed in the aforementioned papers, there are several decision models in the confidence family. Here we use the maximin-EU version, according to which preferences are represented by:

$$V(f) = \min_{p \in D(f)} \mathbb{E}_p u(f(s)) \tag{3}$$

where u is a non-constant affine utility function, Ξ is a closed, convex, min-closed confidence ranking Ξ and D is a function from \mathcal{A} to Ξ , satisfying the following *richness* condition: for every $f \in \mathcal{A} \setminus \mathcal{A}^c$ and $\mathcal{C} \in \Xi \setminus \{\min \Xi\}$, there exist $d \in \mathcal{A}^c$ and $\alpha \in (0, 1]$ such that $D(f_\alpha d) = \mathcal{C}$. This function, called the *cautiousness coefficient* for Ξ , captures the decision maker's ambiguity attitudes, or attitudes to choosing on the basis of limited confidence. We refer to the cited papers for discussion, details and other models in the confidence family, to which the approach developed here can be extended.

When (3) holds for preferences \succeq , we say that the triple (Ξ, D, u) *represents* \succeq . Whenever \succeq is non-degenerate, there is a unique triple (up to positive affine transformation of the utility function) representing \succeq (Hill, 2013), which we refer to as *the representation of \succeq* . We adopt the convention that \succeq is degenerate if and only if it is represented by $\Xi = \{\emptyset\}$ and D the only function from \mathcal{A} to $\{\emptyset\}$. For mere technical convenience, we will suppose throughout that utility is unbounded: $u(\mathcal{X}) = \mathbb{R}$.

We assume that all preferences, ex ante and ex post, are represented according the confidence model (3), and focus on non-degenerate ex ante preferences.

Assumption 1. \succeq and \succeq_E are represented according to (3) for all $E \in \Sigma$, and \succeq is non-degenerate.

Behavioural foundations for this version of the confidence model have been provided in Hill (2013). They can be used to provide a reformulation of this assumption in terms of preferences.

¹⁰ By convention, if Ξ is empty or the family consisting of the empty set, then Ξ^{\min} is taken to be $\{\emptyset\}$.

¹¹ Since the focus is on the update rule, we work with a fairly general form of confidence ranking throughout. Some convenient parametrised special cases are discussed in Hill (2013, §4; 2019, §5).

2.3. Tastes and stakes

A central idea behind the confidence model is that the beliefs one relies on to decide are those held with a level of confidence that is appropriate given the importance of the decision (Hill, 2013, 2019). In the light of this, when fewer beliefs are invoked—i.e. when a larger set of priors is used, say $D(f) \supset D(g)$ —then this is an indication that the decision maker considers the choice of f to be more important than the choice of g : it involves higher stakes. The converse is not necessarily true: a decision maker may use the same beliefs—the same $\mathcal{C} \in \Xi$ —for decisions of differing importance. Indeed, the standard maximin-EU decision rule with a single set of priors (Gilboa and Schmeidler, 1989) is a special case of (3) of just this sort: the same set of probability measures are used for all acts, no matter the stakes involved.

Throughout, we assume that only beliefs—in the context of this model, the confidence ranking Ξ —change on learning, and in particular that there is no change in the utility function or in the level of stakes that a decision is considered to involve.

Assumption 2. For representations (Ξ, D, u) and $\{(\Xi_E, D_E, u_E)\}_{E \in \Sigma}$ of \succeq and $\{\succeq_E\}_{E \in \Sigma}$ respectively:

1. u and u_E are identical up to positive affine transformation for every $E \in \Sigma$;
2. there exists a complete transitive relation \succeq on \mathcal{A} such that for all $E \in \Sigma$ and all $f, g \in \mathcal{A}$, $f \succeq g$ implies $D(f) \supseteq D(g)$ and $D_E(f) \supseteq D_E(g)$.

The first part of this assumption is standard. The second clause states that there is a single notion of stakes (captured by \succeq) that all preferences, ex ante and ex post, can be thought of as respecting. It reflects the assumption that the decision maker’s view of the relative importance of decisions remains constant under learning.

Whilst stated on the models for ease, Assumption 2 can be reformulated in behavioural terms. The first clause corresponds to the standard axiom that preferences over constant acts are unaffected by learning. The latter is built into axiomatisations of the confidence model assuming an exogenously given notion of stakes (Hill, 2013); framework-specific axioms characterise it in setups where stakes are endogenous (Hill, 2015).

Given Assumption 2, \mathcal{A} can be partitioned into *stakes levels* according to \succeq . We use σ_f to denote the stakes level of f : that is, the set of acts having the same stakes as f , $\sigma_f = \{g \in \mathcal{A} : g \succeq f \ \& \ g \preceq f\}$. We use σ, σ' as notation for stakes levels. With this notation, $f \in \sigma$ if f involves stakes of level σ . The obvious order on stakes levels is defined as standard: for stakes levels σ, σ' , $\sigma \geq \sigma'$ if and only if, for all $f \in \sigma$, $f' \in \sigma'$, $f \succeq f'$.

Finally, given a preference relation \succeq represented according to (3) and a stakes level σ , we define the derived relation \succeq^σ as follows: for all $f, g \in \mathcal{A}$, $f \succeq^\sigma g$ if and only if there exists $c, d, d' \in \mathcal{A}^c$ and $\alpha, \alpha' \in (0, 1]$ such that $D(f_\alpha d) = D(g_{\alpha'} d') = D(h)$ for all $h \in \sigma$, $f_\alpha d \succeq c_\alpha d$ and $c_{\alpha'} d' \succeq g_{\alpha'} d'$.¹² As discussed in Hill (2013), $f \succeq^\sigma g$ essentially says that, if the acts were evaluated ‘as if’ they both involved stakes of level σ , then f would be preferred. For example, consider a bet f on the Democrat candidate winning the 2024 US President election, yielding \$1 million if you win and a loss of \$1 million if not, and a similar bet g on the 2028 election, with stakes (winnings and losses) 100000 times less in utility terms. An agent with beliefs that are

¹² This is well-defined because of the richness of D .

more precise and slightly more favourable for the 2024 election might nevertheless prefer g to f because of the difference in stakes: with lower stakes, he can rely on low-confidence beliefs in evaluating g , but not for f . However, if both options were evaluated at the same stakes level—for instance, if bets on the 2024 and 2028 elections with stakes of \$1 million were compared—then f would typically be preferred: i.e. $f \succeq^\sigma g$, where σ is the appropriate stakes level. When $f \succeq^\sigma g$, we say that f is preferred to g at stakes level σ .

2.4. Update

We now formally present the updates rules that we will consider. We shall say that a correspondence $\gamma : X \rightrightarrows Y$ between two ordered sets (X, \geq_X) , (Y, \geq_Y) is *increasing* (resp. *decreasing*) if, for every $y, y' \in Y$, $x, x' \in X$, if $y \in \gamma(x)$, $y' \in \gamma(x')$ and $x \geq_X x'$, then $y \geq_Y y'$ (resp. $y \leq_Y y'$).¹³ We use the natural order, given by containment, on confidence rankings.

Our benchmark update rule is the following.

Definition 1. For confidence rankings Ξ and Ξ_E and an event $E \in \Sigma$, Ξ_E is a (general) confidence update of Ξ by E if there exists a confidence ranking Ξ_{UpdE} and an increasing correspondence $c_{UpdE} : \Xi \rightrightarrows \Xi_{UpdE}$ such that

$$\Xi_E = \left\{ \overline{(\mathcal{C} \cap \mathcal{C}')}_E : \mathcal{C} \in \Xi, \mathcal{C}' \in \Xi_{UpdE} \text{ s.t. } \mathcal{C}' \in c_{UpdE}(\mathcal{C}), \mathcal{C} \cap \mathcal{C}' \neq \emptyset \right\}^{\min} \tag{4}$$

where, for $\mathcal{C} \subseteq \Delta(\Sigma)$ and $E \in \Sigma$, \mathcal{C}_E is the Full Bayesian update defined in (2). The calibration correspondence $c_E : \Xi \rightrightarrows \Xi_E$ is defined by $c_E(\mathcal{C}) = \left\{ \overline{(\mathcal{C} \cap \mathcal{C}')}_E : \mathcal{C}' \in c_{UpdE}(\mathcal{C}), \mathcal{C} \cap \mathcal{C}' \neq \emptyset \right\}$ if there exists $\mathcal{C}' \in c_{UpdE}(\mathcal{C})$ with $\mathcal{C}' \cap \mathcal{C}$ non-empty, and $c_E(\mathcal{C}) = \{\min \Xi_E\}$ otherwise.

This rule generalises the update logic discussed in the Introduction. The information that E is taken to indicate something about how reasonable (prior) probability measures are; however, unlike (1), it is not assumed to amount to a probability threshold at each confidence level. Rather, a set of ‘reasonable’ probability measures in the light of the fact that E has been learnt is specified for each confidence level, representing the conclusions that can be drawn with various levels of confidence. The conclusions are weaker (and the sets are larger) for higher confidence levels, so the information can be represented as confidence ranking, Ξ_{UpdE} . The correspondence c_{UpdE} picks out the appropriate sets in Ξ_{UpdE} for the various confidence levels. Beyond this difference, the rule operates as discussed previously. Whenever the initial beliefs and conclusions drawn from the data at a confidence level are compatible, they are both retained—by taking the intersection of the sets. Whenever they aren’t, neither is retained and the posterior beliefs are inherited from higher confidence levels. Note that if Ξ and Ξ_{UpdE} are closed, convex and min-closed confidence rankings, then Ξ_E defined according to (4) is as well.¹⁴

General confidence update is permissive in how conclusions are drawn from the learnt event E ; we also consider more restrictive special cases.

¹³ A correspondence $\gamma : X \rightrightarrows Y$ is a function from X to $2^Y \setminus \emptyset$. If Y is a lattice, γ is increasing in the defined sense if and only if it is increasing in the Strong Set Order.

¹⁴ Since the Full Bayesian update of a closed set of priors is not necessarily closed (see Section 4.1), the closure is required in (4) to ensure that the ex post confidence ranking is closed. We work with closed confidence rankings for mere convenience (the confidence ranking revealed from preferences is only unique up to closure); the results can be extended to versions of the rule that do not impose closure.

Definition 2. For confidence rankings Ξ and Ξ_E and an event $E \in \Sigma$, Ξ_E is a *probability-threshold confidence update* of Ξ by E if there exists a decreasing correspondence $\rho_E : \Xi \rightrightarrows [0, 1]$ such that

$$\Xi_E = \left\{ \overline{\{p \in \mathcal{C} : p(E) \geq r\}}_E : \mathcal{C} \in \Xi, r \in \rho_E(\mathcal{C}), \{p \in \mathcal{C} : p(E) \geq r\} \neq \emptyset \right\}^{\min} \quad (5)$$

ρ_E is called the *probability-threshold correspondence*.

Probability-threshold confidence update—or *PT-confidence update*, for short—is the special case where the updating confidence ranking Ξ_{UpdE} consists of sets of probability measures satisfying probability thresholds.¹⁵ It is built on the same intuition that learning E indicates something about how reasonable probability measures are: here, those that give too low a probability to E ex ante ‘got it more wrong’ than others, and hence may not be retained at certain confidence levels. At every confidence level, the correspondence ρ_E can be interpreted as providing a threshold that picks out the probability measures retained at that level. If the decision maker has a different set of beliefs at each confidence level, then ρ_E is a function; indeed, modulo some technicalities, the update rule (1) in the Introduction corresponds precisely to this case. However, to accommodate cases where the decision maker holds the same beliefs at different confidence levels—as in the special case of a single set of priors (Section 2.2)—we allow ρ_E to be a correspondence. This allows him to have the same initial beliefs at two different confidence levels, but to consider that observation warrants the use of different probability thresholds.

Given preference relations representable by the confidence model, \succeq and \succeq_E for an event $E \in \Sigma$, we say that \succeq_E is a *general confidence update* of \succeq by E if, for any representations (Ξ, D, u) and (Ξ_E, D_E, u) of \succeq and \succeq_E respectively, Ξ_E is a general confidence update of Ξ and $D_E(f) \in c_E(D(f))$ for all $f \in \mathcal{A}$. PT-confidence update of preferences is defined analogously.

3. Characterising confidence update

We now provide behavioural characterisations of the general and special cases of confidence update.

3.1. General confidence update

A specific comparison of acts with constant acts will play a special role in the axioms below. The preference of an act f over a constant act c betrays that the decision maker values f at least as highly as c . Whilst the value assigned to f may change on learning, the assumption of constant tastes (Assumption 2) ensures that the value of c will not: to that extent, it provides a constant ‘benchmark’. The acts f_{EC} and c coincide whenever E is not the case, in which case the constant benchmark c obtains. A preference for f_{EC} over c thus indicates that, *conditional on E*, f is evaluated as better than the constant benchmark c . This is a special case of the standard definition of conditional preferences under expected utility, which compares $f_E h$ and $c_E h$. (For expected utility, unlike for ambiguity models, this comparison is independent of h .) The specific case used here—where $h = c$ —is the only one where one of the acts is guaranteed

¹⁵ More formally, Ξ_E is a probability-threshold confidence update of Ξ by E if and only if Ξ_E is a general confidence update of Ξ by E , with $\Xi_{UpdE} = \{\{p \in \Delta(\Sigma) : p(E) \geq x\} : x \in [0, 1]\}$ and $c_{UpdE}(\mathcal{C}) = \{\{p \in \Delta(\Sigma) : p(E) \geq r\} : r \in \rho_E(\mathcal{C})\}$ for all $\mathcal{C} \in \Xi$.

to be constant, and hence has a value that is independent of beliefs. To the extent that it ties in with the use of constant acts as a benchmark for evaluating others, preference comparisons of f_{EC} and c thus provide a natural conception of conditional preferences. As discussed in Section 4.2.1, conditional preferences are important, because the central issue in update concerns what happens to conditional beliefs.

We introduce the following terminology. For an event $E \in \Sigma$ and a stakes level σ , we say that σ is *E-resilient* if, for all $f \in \mathcal{A}$, $c \in \mathcal{A}^c$, if $f_{EC} \succeq^\sigma c$, then $f_{EC} \succeq_E^\sigma c$. *E-resilient* stakes levels are those for which all relevant ex ante conditional preferences are retained on learning E : if f is evaluated as better than c conditional on E prior to learning, then it continues to enjoy this evaluation afterwards.

These concepts are familiar in the literature on (Bayesian and non-Bayesian) updating. For instance, Pires's (2002) axiomatisation of Full Bayesian update under the maximin-EU model involves a related notion of conditional preferences, and his main axiom is a strengthening of the condition that all stakes levels are *E-resilient*, for every non-null E . As discussed in Section 4.2.1, Dynamic Consistency also imposes a stronger conditional-preference preservation property than *E-resilience*, for all stakes levels.

3.1.1. Main axiom

The following is the central behavioural axiom behind confidence update.

Axiom (Confidence Consistency). For all stakes levels σ, σ' with $\sigma' > \sigma$ and every non-null $E \in \Sigma$, if σ is *E-resilient*, then so is σ' .

Confidence Consistency translates the maxim mooted in the Introduction: retain those conditional beliefs in which you are more confident, and relinquish those in which you have less confidence. If σ is *E-resilient*, then all ex ante conditional evaluations of acts relative to 'benchmark' constant acts are retained ex post. This indicates that the beliefs underlying these preferences are retained on update. Confidence consistency implies that if all such conditional preferences are retained at some stakes level, then the conditional preferences at any higher stakes level are also retained. If the decision maker is confident enough in the beliefs underlying the former preferences to hold onto them in the face of the information E , then he will also hold onto the beliefs underlying the latter preferences. This is precisely as the maxim demands: if he retains beliefs held at a given level of confidence, then he certainly cannot relinquish beliefs held with higher confidence, for he should have relinquished the former beliefs first!

3.1.2. Other axioms

Now consider the following axioms.

Axiom (Consequentialism). For every non-null $E \in \Sigma$, if $f(s) = g(s)$ for all $s \in E$, then $f \sim_E g$.

Axiom (Non-degeneracy). For every non-null $E \in \Sigma$, \succeq_E is non-degenerate.

Axiom (Information-based Learning). For every $f \in \mathcal{A}$, $c \in \mathcal{A}^c$ and $E \in \Sigma$, if $f \not\prec_E^\sigma c$ for every *E-resilient* stakes level σ , then $f \not\prec_E^{\sigma'} c$ for every σ' .

Consequentialism is a well-known and relatively standard axiom in the dynamic context; see e.g. Epstein and Le Breton (1993); Ghirardato (2002) for further discussion of it. Non-degeneracy is the standard property that update by non-null events yields non-degenerate preferences.

Confidence Consistency concerns what happens when learning E does not shake beliefs held with a certain level of confidence. By contrast, Information-based Learning constrains what happens when learning E does shake beliefs at a particular confidence level—that is, at stakes levels which are not E -resilient, and hence where some ex ante conditional preferences are not retained on learning. The condition basically implies that preferences at these stakes levels are fully determined by preferences at higher, E -resilient stakes levels, where the information can be incorporated without relinquishing ex ante beliefs. Hence it demands that learning is entirely driven by the new information E . If learning E undermines beliefs held only to a low level of confidence, they will not be replaced with anything specific. The information is only understood as saying that such low-confidence beliefs are inappropriate, but not as specifying other beliefs to replace them, except those beliefs inherited from higher confidence levels.

We call these three axioms the *Basic Axioms*.

3.1.3. Representation

Confidence Consistency and the Basic Axioms yield our most general update rule.

Theorem 1. *Let \succeq and $\{\succeq_E\}_{E \in \Sigma}$ satisfy Assumptions 1 and 2. They satisfy Confidence Consistency and the Basic Axioms if and only if, for every non-null $E \in \Sigma$, \succeq_E is a general confidence update of \succeq .*

So, in the presence of the Basic Axioms, Confidence Consistency characterises the heart of the proposed approach, namely the general confidence update rule. In fact, the central behavioural properties of the approach essentially boil down to Confidence Consistency, Consequentialism and Non-degeneracy. Information-based Learning merely controls what happens at the bottom of the confidence ranking, where there is incompatibility with prior beliefs. It can be shown that in its absence, the essence of confidence update is retained, except at confidence levels at the bottom of the ranking.

3.2. Probability-threshold confidence update

3.2.1. Axioms

To obtain the specification of general confidence update involving probability thresholds in Definition 2, consider the following axioms.

Axiom (Probability Consistency). *Consider any non-null $E \in \Sigma$ and E -resilient stakes levels $\sigma \leq \sigma'$. For every $\lambda \in (0, 1]$ and $f, g \in \mathcal{A}$, if, for every $c, \underline{c} \in \mathcal{A}^c$ with $c \succ \underline{c}$, $f_{EC} \sim^{\sigma'} \underline{c}_\lambda c$ implies $(f_{EC})_{\frac{1}{2}}(c_E \underline{c}) \succeq_E^{\sigma'} c_{\frac{1}{2}} \underline{c}$, then, for every $d, \underline{d} \in \mathcal{A}^c$ with $d \succ \underline{d}$, $g_{Ed} \sim^\sigma \underline{d}_\lambda d$ implies $(g_{Ed})_{\frac{1}{2}}(d_E \underline{d}) \succeq_E^\sigma d_{\frac{1}{2}} \underline{d}$.*

Axiom (Null consistency). *For every non-null $E \in \Sigma$, E -resilient stakes level σ , $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, if $f_{Ee} \leq^\sigma c$ for all $e \in \mathcal{A}^c$, then $f_{EC} \leq_E^\sigma c$.*

To interpret Probability Consistency, note that $(f_{EC})_{\frac{1}{2}}(c_E \underline{c})$ is a 50-50 mixture of f_{EC} with a bet on the event E —the act $c_E \underline{c}$ —whereas $c_{\frac{1}{2}} \underline{c}$ is a 50-50 mixture of $\underline{c}_\lambda c$ with a bet yielding the

winning option c with probability λ —that is, $c_\lambda c$. So if a decision maker weakly prefers the first bet (f_{EC} ; the bet on E) over the second ($c_\lambda c$; the bet with probability λ of winning) in each case, then she would weakly prefer the first mixture ($(f_{EC})_{\frac{1}{2}}(c_{E\bar{C}})$) over the second ($(c_{\frac{1}{2}}c)$).¹⁶ This can be thought of as an ‘implication’ of the previous two preferences. But a weak preference for the bet on E over that with probability λ (i.e. $c_{E\bar{C}} \succeq c_\lambda c$) betrays a judgement that the probability of E is λ or greater. So a weak preference for $(f_{EC})_{\frac{1}{2}}(c_{E\bar{C}})$ over $c_{\frac{1}{2}}c$ would be an ‘implication’ of a prior weak preference for f_{EC} over $c_\lambda c$ and a judgement that the probability of E is λ or greater. In the light of this, the axiom says that if the decision maker’s ex post preferences at some stakes level are consistent with all such ‘implications’ of the judgement that the probability of E is λ or greater (i.e. she weakly prefers each relevant $(f_{EC})_{\frac{1}{2}}(c_{E\bar{C}})$ ex post), then they remain consistent with all such ‘implications’ of that judgement at any lower E -resilient stakes level. In other words, if the decision maker’s preferences are consistent with her incorporating the opinion that E is more probable than λ at some stakes level, then they are consistent with her incorporating that opinion at any lower stakes level. This is the sort of pattern one would expect given the intuitions about the information purveyed on learning E : if at some confidence level, the decision maker considers the observation that E to warrant a judgement that its probability was greater than λ , then she still considers it to warrant that judgement at any lower confidence level.

Null consistency concerns the case where the ex ante evaluation of an act f_{Ee} remains bounded by a constant act c , no matter how attractive e is. This indicates that E^c involves a certain form of nullness—certainly, if its probability were bounded away from 0 across the relevant set of priors, then such preferences would not occur. The axiom says that, in such cases of ex ante nullness, the ex post evaluation concerning f remains bounded by c . This is reasonable: if the event E was already treated as if it had probability 1 in that region ex ante, then on learning E , the decision maker’s evaluation of f cannot rise much.

3.2.2. Representation

Theorem 2. *Let \succeq , $\{\succeq_E\}_{E \in \Sigma}$ satisfy Assumptions 1 and 2. Then they satisfy Confidence Consistency, the Basic Axioms, Probability Consistency and Null consistency if and only if, for every non-null $E \in \Sigma$, \succeq_E is a probability-threshold confidence update of \succeq .*

In the presence of Null consistency, Probability Consistency thus characterises the probability-threshold specification of general confidence update discussed in the Introduction and Section 2.4. In other words, it guarantees the existence of a probability-threshold correspondence that characterises update according to (5). The following result characterises the uniqueness of this correspondence.

Proposition 1. *Let \succeq , $\{\succeq_E\}_{E \in \Sigma}$ satisfy the conditions in Theorem 2, with the former represented by (Ξ, D, u) . There exists a unique maximal probability-threshold correspondence $\rho_E : \Xi \rightarrow [0, 1]$ representing the update of Ξ by E , in the following sense: for every other ρ'_E representing the update by E and for every $C \in \Xi$, if $y \in \rho'_E(C)$, then there exists $x \in \rho_E(C)$ with $x \geq y$. Moreover, if for every stakes level σ , there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ such that $f_{EC} \succeq_E^\sigma c$ but*

¹⁶ Since the second bets are both constant acts, this is a consequence of the uncertainty aversion of the maximin-EU model.

$f_{EC} \not\prec^{\sigma} c$, then there exists $C \in \Xi$ such that the correspondence ρ_E representing the update by E is unique on all $C' \in \Xi$ with $C' \supset C$.

There is a unique maximal probability-threshold correspondence, in the sense that it yields values higher than those given by any other correspondence representing the update. Moreover, whenever something is learnt (i.e. preferences change) at every stakes level, then on all sufficiently large confidence levels, the probability-threshold correspondence is uniquely revealed from preferences.

3.3. Further special cases

PT-confidence update (5) involves a probability-threshold correspondence ρ_E for each event E , without assuming any relationship between them. However, the approach can easily accommodate richer structures, involving closer relationships between the probability-threshold correspondences for different events. These may be useful for applications, or in connecting the approach to existing work in statistics (Section 5). By way of illustration, we provide an axiomatisation of the simplest such special case: where a single probability-threshold correspondence represents update for all (non-null) events.¹⁷ To this end, consider the following strengthening of Probability Consistency.

Axiom (Strong Probability Consistency). Consider any non-null $E, F \in \Sigma$ and E - and F -resilient stakes levels $\sigma \leq \sigma'$. For every $\lambda \in (0, 1]$ and $f, g \in \mathcal{A}$, if, for every $c, \underline{c} \in \mathcal{A}^c$ with $c \succ \underline{c}$, $f_{EC} \sim^{\sigma'} \underline{c}_{\lambda} c$ implies $(f_{EC})_{\frac{1}{2}}(c_{E\underline{c}}) \succeq_E^{\sigma'} c_{\frac{1}{2}} \underline{c}$, then, for every $d, \underline{d} \in \mathcal{A}^c$ with $d \succ \underline{d}$, $g_{Fd} \sim^{\sigma} \underline{d}_{\lambda} d$ implies $(g_{Fd})_{\frac{1}{2}}(d_{F\underline{d}}) \succeq_F^{\sigma} d_{\frac{1}{2}} \underline{d}$.

The central difference in this axiom with respect to Probability Consistency is that it compares across different events; apart from that, the interpretation in terms of the lower stakes levels retaining the judgements whose ‘implications’ are respected at higher stakes levels remains the same. This strengthening yields the desired special case.

Proposition 2. Let $\succeq, \{\succeq_E\}_{E \in \Sigma}$ satisfy Assumptions 1 and 2. Then they satisfy Confidence Consistency, the Basic Axioms, Null consistency and Strong Probability Consistency if and only if there exists a probability-threshold correspondence $\rho : \Xi \rightrightarrows [0, 1]$ such that, for every non-null $E \in \Sigma$, \succeq_E is a confidence update of \succeq by E represented by ρ . The uniqueness of ρ is as in Proposition 1.

3.4. Comparative statics

In this section, we take a brief look at the comparative statics of the PT-confidence update rule, as concerns ex post ambiguity aversion. We adopt a standard definition of comparative

¹⁷ The aim of this exercise is to illustrate the strength of the approach; we by no means wish to suggest that equality of probability-threshold correspondences is reasonable or desirable. For instance, it seems more reasonable to look at equal likelihood ratio-thresholds across events; however, note that in our extremely general framework (where the space of probability measures is the whole of $\Delta(\Sigma)$), the likelihood ratio coincides with the likelihood. Restricting the space of measures considered and using an adapted version of the techniques presented here can provide a likelihood ratio-version of the result below; details go beyond the scope of the current paper.

ambiguity aversion (Ghirardato and Marinacci, 2002), according to which decision maker \succeq' is more ambiguity averse than \succeq if and only if, for all $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, if $f \succeq' c$, then $f \succeq c$.

Proposition 3. *Let $\succeq, \{\succeq_E\}_{E \in \Sigma}$ and $\succeq', \{\succeq'_E\}_{E \in \Sigma}$ be two families satisfying Assumptions 1 and 2, and the conditions in Theorem 2, and suppose that $\succeq = \succeq'$. Let (Ξ, D, u) be the representation of \succeq , and let $\{\rho_E\}$ and $\{\rho'_E\}$ be the families of maximal correspondences as specified in Proposition 1 representing updates yielding $\{\succeq_E\}_{E \in \Sigma}$ and $\{\succeq'_E\}_{E \in \Sigma}$ respectively. Then the following are equivalent, for each non-null event E :*

- (i) *for every E -resilient stakes level σ according to \succeq , \succeq'^σ_E is more ambiguity averse than \succeq_E ;*
- (ii) *for every $C \in \Xi$ and $y \in \rho'_E(C)$, there exists $x \in \rho_E(C)$ with $x \geq y$.*

This result sheds light on the role of the probability-threshold correspondence. For decision makers with identical ex ante preferences, differences in ex post ambiguity attitude at stakes levels where the relevant conditional preferences are retained on learning correspond to differences in the probability-threshold correspondence. The latter essentially reflects the strength of the conclusions a decision maker is willing to draw from a given observation for each confidence level: as discussed in the Introduction, it reflects how ‘wrong’ a probability measure has to be ex ante about the new information for it to be ruled out as plausible. The higher the probability-threshold at a given confidence level, the stricter this constraint, and hence the stronger the implicit conclusions the decision maker is drawing from the data. So, if one decision maker always uses a higher probability threshold than another, the former can be thought of as more daring, or less cautious, in the conclusions he is prepared to draw from the same data. This translates to him being less ambiguity averse ex post.

As discussed below (Section 5), the probability-threshold correspondence plays a similar role to significance levels in statistics, with the difference that it assigns a significance level to each level of confidence. Decision makers which differ in the probability-threshold correspondence (or, equivalently, *ceteris paribus*, ex post ambiguity aversion) can thus be thought of, roughly, as differing in the significance level they deem appropriate for a given level of confidence.

4. Situating confidence update

As noted in the Introduction, update of ambiguous beliefs presents a certain number of challenges, concerning in particular complete ignorance and surprising events. We shall now consider how the proposed approach fares with respect to these challenges, and compares to existing update rules in the literature.

4.1. Complete ignorance, and other updating rules for ambiguous beliefs

The most notable generally-applicable consequentialist update rules that have been proposed and axiomatised for multiple prior models, and in particular the maximin-EU model (Gilboa and Schmeidler, 1989), are Full Bayesian (Pires, 2002; Walley, 1991) and Maximum Likelihood update (Gilboa and Schmeidler, 1993; Dempster, 1967).¹⁸ As Gilboa and Marinacci (2013) point

¹⁸ Other approaches to update in the literature drop consequentialism (Hanany and Klibanoff, 2007), restrict to ex ante beliefs satisfying a particular property with respect to a given filtration of events representing the potential new

out, both are extreme, which of course sheds doubt on their descriptive adequacy as well as their normative validity. Some other, apparently milder, rules have been proposed, for instance in Epstein and Schneider (2007). As we shall now show, all of these rules come out as special instances of confidence update.

Let us assume that the initial confidence ranking is a singleton containing the closed convex set of probability measures $\mathcal{P} \subseteq \Delta(\Sigma)$, so initial preferences are maximin-EU (Gilboa and Schmeidler, 1989). For update by an event E , applying the PT-confidence update rule to these preferences yields preferences represented according to (3), with confidence ranking

$$\Xi_E = \left\{ \overline{\{p \in \mathcal{P} : p(E) \geq r\}}_E : r \in R_E \text{ s.t. } \{p \in \mathcal{P} : p(E) \geq r\} \neq \emptyset \right\}^{min} \tag{6}$$

where $\rho_E(\mathcal{P}) = R_E \subseteq [0, 1]$. As discussed previously, the confidence rule allows one to distinguish on update according to how reasonable probability measures are in the light of the information. Hence it can yield, even for a degenerate initial confidence ranking (i.e. a single set of probability measures) a richer posterior confidence ranking; indeed, this will typically be the case whenever R_E is not a singleton. We shall say that a PT-confidence update by E is *maximally refined* whenever $R_E = [0, 1]$, in which case (6) simply becomes:

$$\Xi_E^{mr} = \left\{ \overline{\{p \in \mathcal{P} : p(E) \geq r\}}_E : r \in [0, 1] \text{ s.t. } \{p \in \mathcal{P} : p(E) \geq r\} \neq \emptyset \right\} \tag{7}$$

The aforementioned update rules from the literature correspond to particular confidence levels in the confidence ranking resulting from the maximally refined application of the PT-confidence update rule. Consider, for instance, the largest set in Ξ_E^{mr} , capturing the beliefs held with the highest level of confidence. This is (the closure of) \mathcal{P}_E : the Full Bayesian update of the initial set of priors \mathcal{P} (see (2)). On the other hand, the smallest set in Ξ_E^{mr} —capturing all beliefs held, no matter how little confidence there is in them—is clearly (the closure of) $\{p \in \mathcal{P} : p(E) = \max_{q \in \mathcal{P}} q(E)\}_E$: the Maximum-Likelihood update of \mathcal{P} . Moreover, for every ‘significance level’ $\alpha \in [0, 1]$, the ‘classical-style’ update which retains all probability measures giving a probability greater than α to E —i.e. $\{p \in \mathcal{P} : p(E) \geq \alpha\}_E$ —coincides (up to closure) with some non-extremal set in Ξ_E^{mr} , corresponding to some intermediate confidence level. Whilst axiomatisations of the two previous rules are well known in the economics literature (Gilboa and Schmeidler, 1993; Pires, 2002), the results in Section 3 also yield a behavioural characterisation of this last rule.¹⁹ Moreover, they provide what to our knowledge is the first unified behavioural characterisation of this whole family of rules.

We illustrate these points, and the consequences for ‘complete ignorance’ cases, on a simple version of the venture capitalist example from the Introduction.

Example 1 (Complete Ignorance). Recall that the venture capitalist is considering an entirely new drug about which she knows absolutely nothing. This can be cast in a statistical decision-style framework as follows. For patient n having the illness and treated with the drug, the state space $S_n = S = \{s, f\}$ (s for success in treating the illness, f for failure); we can take the full

information (Epstein and Schneider, 2003), or define update in a multi-stage context with a given information structure, with update depending on the first-stage partition (Gul and Pesendorfer, 2021). Note that in the sort of challenging cases mentioned above—complete ignorance, learning on surprising or null events—a decision maker would not typically have a full and correct conception of the information structure (filtration) she is facing.

¹⁹ After completing the paper, our attention was drawn to Kovach (2015), which develops a different, specific axiomatisation of the last rule.

state space $\mathcal{S} = S^\infty$, with the product σ -algebra Σ . The standard statistical decision framework assumes in addition a parameter space Θ . Under the assumption that patient trials are IID, we can take $\Theta = [0, 1]$, with each $\theta \in \Theta$ associated to the probability distribution $\ell(\bullet/\theta)$ over the S (i.e. an element of $\Delta(S)$),²⁰ where the probability of success $\ell(s/\theta) = \theta$. Each $\theta \in \Theta$ thus generates the distribution $\ell^\infty(\bullet/\theta)$ over \mathcal{S} . Just as a distribution over the parameter space $\mu \in \Delta(\Theta)$ generates a *predictive* distribution over the full state space $\bar{\mu} = \int_{\Theta} \ell^\infty(\bullet/\theta) d\mu(\theta) \in \Delta(\mathcal{S})$, a set of probability measures $\mathcal{M} \subseteq \Delta(\Theta)$ generates a set of probability measures $\overline{\mathcal{M}} \subseteq \Delta(\mathcal{S})$, defined as follows:

$$\overline{\mathcal{M}} = \left\{ \int_{\Theta} \ell^\infty(\bullet/\theta) d\mu(\theta) : \mu \in \mathcal{M} \right\} \tag{8}$$

We adopt a multiple prior representation, \mathcal{M} , of ex ante beliefs. Following the standard way of representing a complete lack of prior knowledge about the success of the drug in this context, we set $\mathcal{M} = \Delta(\Theta)$. In particular, \mathcal{M} contains every Dirac measure; we denote by μ_θ the Dirac measure putting all weight on θ .

Suppose that the venture capitalist observes 100 patient trials, 75 of which were successes for the drug—call this event t_{100} . We consider her posterior belief concerning s_{101} —success of the drug on the next patient. Applying the maximally refined confidence update, as in (7), to $\overline{\mathcal{M}}$ yields:

$$\begin{aligned} \Xi_{t_{100}} = & \left\{ \overline{\left\{ \int_{\Theta} \ell^\infty(\bullet/\theta) d\mu_{t_{100}}(\theta) : \mu \in \Delta(\Theta), \int_{\Theta} \ell^\infty(t_{100}/\theta) d\mu(\theta) \geq r, \int_{\Theta} \ell^\infty(t_{100}/\theta) d\mu(\theta) > 0 \right\}} \right. \\ & \left. : r \in \left[0, \max_{\mu \in \Delta(\Theta)} \int_{\Theta} \ell^\infty(t_{100}/\theta) d\mu(\theta) \right] \right\} \end{aligned} \tag{9}$$

where, as standard,

$$\mu_{t_{100}}(A) = \frac{\int_A \ell^\infty(t_{100}/\theta) d\mu(\theta)}{\int_{\Theta} \ell^\infty(t_{100}/\theta) d\mu(\theta)} \tag{10}$$

for any (measurable) $A \subseteq \Theta$.

The Full Bayesian update of $\overline{\mathcal{M}}$ on t_{100} is $\left\{ \int_{\Theta} \ell^\infty(\bullet/\theta) d\mu_{t_{100}}(\theta) : \mu \in \Delta(\Theta), \int_{\Theta} \ell^\infty(t_{100}/\theta) d\mu(\theta) > 0 \right\}$. Up to closure, this coincides with the set of probability measures corresponding to the highest confidence level in $\Xi_{t_{100}}$, $\max \Xi_{t_{100}}$. In particular, since the only $\mu \in \Delta(\Theta)$ with $\int_{\Theta} \ell^\infty(t_{100}/\theta) d\mu(\theta) = 0$ are the Dirac measures μ_0 and μ_1 , this set is $\overline{\mathcal{M}}$ itself. Since, under the maximin-EU rule, a set of priors is behaviourally indistinguishable from its closure—both yield the same preferences—this means that, under Full Bayesian update, the decision maker’s preferences do not change on learning. Full Bayesian update thus allows for no learning in such cases of ex ante complete ignorance. This property, known as the issue of ‘complete ignorance’ or ‘vacuous priors’ (Walley, 1991, §§6.6.1, 9.3), has been the topic of intense debate in some

²⁰ We use $\Delta(S)$ in the context of this example to denote the set of probability distributions over S , and similarly for $\Delta(\mathcal{S})$ and $\Delta(\Theta)$.

circles, where it is considered a major challenge for non-Bayesian accounts (e.g. Joyce, 2010; Bradley, 2017; Vallinder, 2017).

On the other hand, Maximum Likelihood update yields $\{\int_{\Theta} \ell^{\infty}(\bullet/\theta)d\mu_{t_{100}}(\theta) : \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta) = \max_{\mu \in \Delta(\Theta)} \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta)\}$, which coincides up to closure with the set of probability measures corresponding to the lowest confidence level in $\Xi_{t_{100}}$, namely $\min \Xi_{t_{100}}$. Since $\ell^{\infty}(t_{100}/0.75) > \ell^{\infty}(t_{100}/\theta)$ for every $\theta \neq 0.75$, this set is the singleton containing $\ell^{\infty}(\bullet/0.75)$. After learning t_{100} , the decision maker using this rule thus assigns a precise probability of 0.75 to s_{101} . So the Maximum Likelihood update rule goes to the opposite extreme: the decision maker settles on a precise opinion after a finite number of observations, and indeed, does so even if the number of observations is very small.

More reasonable than these extremes are the update rules one gets when restricting to intermediate confidence levels. Up to closure, these yield posterior sets of probability measures such as $C^{\alpha} = \{\int_{\Theta} \ell^{\infty}(\bullet/\theta)d\mu_{t_{100}}(\theta) : \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta) \geq \alpha\}$, for $\alpha \in [0, \max_{\mu \in \Delta(\Theta)} \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta)]$. For non-extreme α , these sets are neither as imprecise as \overline{M} , nor as specific as a singleton. As a simple illustration, suppose the initial set of priors is the set of Dirac measures, $M^D = \{\mu_{\theta} : \theta \in \Theta\}$. In this case, the maximally refined confidence update of \overline{M}^D yields²¹:

$$\Xi_{t_{100}}^D = \left\{ \overline{\{\ell^{\infty}(\bullet/\theta) : \theta \in \Theta, \ell^{\infty}(t_{100}/\theta) \geq r, \ell^{\infty}(t_{100}/\theta) > 0\}} : r \in [0, 1] \right\} \quad (11)$$

which, at intermediate confidence levels, involves sets of the form $\{\ell^{\infty}(\bullet/\theta) : \theta \in \Theta, \ell^{\infty}(t_{100}/\theta) \geq \alpha\}$, up to closure. (Full Bayesian and Maximum Likelihood update on this set yields analogous results to those above.) It is clear that these sets are smaller for larger values of α , which correspond in turn to lower confidence levels. Furthermore, they will typically be non-extremal.

Moreover, setting $\beta = \frac{\alpha}{\max_{\mu \in M} \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta)}$, we can rewrite $C^{\alpha} = \{\int_{\Theta} \ell^{\infty}(\bullet/\theta)d\mu_{t_{100}}(\theta) : \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta) \geq \beta \max_{\mu \in M} \int_{\Theta} \ell^{\infty}(t_{100}/\theta)d\mu(\theta)\}$ for $\beta \in [0, 1]$. This is the essence of the update rule proposed by Epstein and Schneider (2007, Eqn (6)), albeit in a recursive setup involving incomplete learning.²² As noted previously, it falls out as a consequence of confidence update.

On the experimental front, there is relatively little research to date on the update of multiple priors in situations of complete ignorance. Cohen et al. (2000) use a Ellsberg-style binary choice task to compare Full Bayesian against Multiple Likelihood, finding more evidence for the former. Baillon et al. (2018) elicit indices they connect with multiple priors in an experiment involving differing amounts of information on past stock values when predicting future ones, finding that the ‘size of the set of priors’ decreases on learning, without reaching a singleton. This finding is inconsistent with the Maximum Likelihood rule. Bland and Rosokha (2019), using a mixture model approach on a bag-and-balls experiment, find evidence for the moderate Epstein and Schneider (2007) update among non-Bayesian subjects, but little for the more extreme rules. Since these studies elicit preferences using small-to-medium-stakes bets, for which intermediate confidence levels are most relevant, the confidence approach is consistent with all of their findings. We know of no experimental evidence at present pertaining to its further predictions, discussed below, about the relationship between update and stakes in the ex post decision, and suggest this as an area for future experimental research.

²¹ Recall (Section 2.2) that we do not restrict to convex confidence rankings in this section. See also Section 5.

²² Although we have illustrated the relationship with their update rule in a standard IID context, it is possible to write their incomplete learning model in our general framework and recover their version of the rule.

Confidence update thus offers a new perspective on existing update rules for ambiguous beliefs, and with it a new resolution of the ‘complete ignorance’ problem. Full Bayesian update is what you get if you use the maximally refined version of the confidence update rule, but then only restrict to beliefs in which you have most confidence. It thus basically retains only conclusions that can be gleaned from observation with maximal confidence, ignoring the rest. As such, it comes to appear as a particularly cautious update rule. Maximum-likelihood update, on the other hand, is what you get if you apply confidence update and then allow yourself to rely on all beliefs, even those held with little confidence. It thus admits any conclusion that can be gleaned from observation, no matter the confidence with which it can be deduced: it makes the boldest use possible of observations. The third sort of rule described above corresponds to taking beliefs held to an intermediate level of confidence, and hence embodies an intermediate level of caution. Note that this latter rule is perhaps the closest to much practice: taking the set corresponding to a probability threshold of 0.01, for instance, would be consistent with the classical practice in statistics of taking a 1% significance value (see also Section 5).

These ‘standard’ rules thus turn out to differ not in the underlying update mechanism: they are all retrievable from a single PT-confidence update. Rather, they differ in the confidence level that they embrace in posterior beliefs. However, under the confidence approach, the appropriate confidence level for an ex post decision depends on the importance of the decision and the cautiousness coefficient—which, recall, reflects ambiguity attitude (Section 2.2; see Hill, 2013, 2019 for extended discussion). This new perspective thus suggests that the aforementioned rules are in fact confounding update with ex post ambiguity attitude. Full Bayesian update, for instance, is so cautious because it implicitly recommends, even for the most trivial decision, demanding maximal confidence in the beliefs used—and this is at the heart of its problems with complete ignorance cases. By contrast, the confidence approach does not require one to settle on a *single* confidence level for *all* updates and decisions. Different ex post decisions and different ambiguity attitudes will call for different confidence levels. Whilst in very high stakes decisions, a decision maker may behave as if he is using Full Bayesian update, when the stakes are lower, his ex post choices will be characterised by less extreme interim rules. So whilst, in the troublesome ‘complete ignorance’ cases, the confidence approach recognises that decision makers may behave as if there was no learning when the stakes are extremely high or as if they learn very fast when the stakes are very low, in medium-stakes decisions, behaviour will be consistent with non-trivial yet moderate learning.²³ Confidence update thus provides a generalisation of standard approaches that can situate and resolve the tension between them, and solve the ‘complete ignorance’ problem mentioned in the Introduction.

4.2. Conditional beliefs and surprising or null events

We now turn to the consequences of the confidence approach for updating on surprising or null events. On this front, standard Bayesian update serves as the natural benchmark, since all the multiple prior update rules discussed in the previous section coincide with it when prior beliefs are precise probabilities. We thus start by comparing confidence update with Bayesian conditionalisation, notably on how they deal with conditional beliefs.

²³ Moreover, such a decision maker will behave as if the ex ante ‘complete ignorance’ set of priors contracts more for less important ex post decisions, which require less confidence and hence admit drawing bolder conclusions from observation.

4.2.1. Bayesian conditionalisation and Dynamic Consistency

It is well-known that Bayesian conditionalisation relies on the assumption that conditional probabilities on a non-null event E are unchanged after learning E (e.g. Jeffrey, 1992; Bradley, 2005; Dietrich et al., 2016). Translated into confidence terms, the assumption that the conditional probability of an event F given E is unaffected by learning E essentially boils down to the assumption that the decision maker has sufficient confidence in his judgement about the conditional probability of F given E to retain it in the face of the new information. However, his confidence is a fact about his ex ante beliefs, as encapsulated in his confidence ranking. Indeed, it is a straightforward consequence of confidence update that whenever he is maximally confident in his judgement about the conditional probability of F given E , his conditional beliefs with respect to these events will be invariant. We summarize this in the following fact.

Fact 1. *Let Ξ be a confidence ranking with $\min \Xi = \{p\}$ and $E, F \in \Sigma$ with E non-null. If $q(F/E) = p(F/E)$ for all $q \in \max \Xi$, then for any Ξ_E resulting from a (general) confidence update of Ξ by E , $q'(F/E) = p(F/E)$ for all $q' \in \max \Xi_E$. In particular, if $\min \Xi_E$ is a singleton containing p_E , then $p_E(F/E) = p(F/E)$.*

Confidence rankings containing a singleton set are discussed and characterised in Hill (2013), where they were called *centred* confidence rankings. Decision makers represented by such confidence rankings are Bayesians with confidence: they can assign a precise probability value to any event, but may have limited confidence in some of these assignments. They are thus a natural context for exploring the relationship with standard Bayesian techniques.

The previous observation suggests that the essence of Bayesian update boils down to a property of the decision maker's ex ante beliefs: namely, a large amount of confidence in conditional probabilities. Indeed, it is straightforward to check that if the decision maker is maximally confident in all conditional probabilities, then we return to the Bayesian special case: the confidence ranking contains only one set, which is a singleton. So the proposal here diverges from Bayesianism insofar as it acknowledges that decision makers might, not unreasonably, have limited confidence in some of their conditional probability judgements. In the face of certain information, they may thus relinquish some conditional probability judgements in order to retain others—and hence violate the central tenet behind Bayesian conditionalisation.

Dynamic Consistency represents the behavioural counterpart of the aforementioned invariance property: in the presence of other basic axioms, it is equivalent to the statement that preferences conditional on E are invariant on learning E (Ghirardato, 2002, Lemma 1). On the behavioural front, confidence update leads to an analogous weakening of Dynamic Consistency: the Confidence Consistency axiom allows relinquishing some conditional preferences, whereas Dynamic Consistency preserves them all.²⁴

Whilst this is not the place for an extended discussion of Dynamic Consistency's normative credentials, note that its defense is strongest in cases where decision makers have full and correct ex ante knowledge of the information structure they are faced with (Hill, 2020). In such cases, decision makers will typically be very confident in their conditional beliefs, and hence confidence

²⁴ More precisely, using the terminology introduced in Section 2: Dynamic Consistency implies that $f_E c \succeq^\sigma c$ if and only if $f_E c \succeq_E^\sigma c$ for all f, c and stakes levels σ (and is equivalent to this condition in our setup whenever ex ante preferences are expected utility), whereas Confidence Consistency clearly weakens this condition.

update will coincide with Bayesian conditionalisation, and satisfy Dynamic Consistency (Fact 1). By contrast, learning surprising or null events often give decision makers reason to question their prior conceptions; the assumption of full and correct ex ante awareness underpinning Dynamic Consistency is ill-adapted to such cases. Indeed, as noted in the Introduction, updates on surprising or null events are particularly challenging for Bayesian update. By eschewing Bayesianism’s insistence on the invariance of conditional beliefs and preferences, confidence update provides a more constructive treatment of such cases.

4.2.2. *Scientific discovery, crises and surprises*

Consider the following example, again concerning a new drug.

Example 2. Prior to a sequence of successive trials of a drug, a decision maker is fairly sure that the process is IID, though she is unsure of the probability of success. Suppose (to avoid issues with improper priors) that she uses the standard statistical setup for IID processes (see Example 1, from which we borrow notation), with as parameter space $\Theta = \{0, 0.1, 0.2, \dots, 0.9, 1\}$ where $\ell(s/\theta) = \theta$ for each $\theta \in \Theta$. She is Bayesian, and takes a uniform prior μ on Θ , which generates the predictive $\bar{\mu} \in \Delta(\mathcal{S})$. She then observes 10 000 patient trials, which turn out as follows: $(s, f, s, f, \dots, s, f, s, f)$. Let us call this history of observed trials t^{10000} .

Under Bayesian conditionalisation, her posterior probability for the next trial being a success would coincide with her prior conditional probability, taking the value 0.5 ($p_{t^{10000}}(s_{10001}) = p_{t^{10000}}(s_{10001}/t^{10000}) = \bar{\mu}(s_{10001}/t^{10000}) = 0.5$). Similarly, Dynamic Consistency insists that her ex post evaluation of bets on the 10 001st trial should coincide with her ex ante conditional evaluation: so she is indifferent between betting on success or failure in both cases. However, whilst under the IID assumption the sequence t^{10000} is as probable as every other sequence involving 5 000 successes out of 10 000 trials, the particular pattern in fact seems rather surprising, and may give the decision maker reasonable grounds to question the IID nature of the process. If so, she would tend to think that a success on the next trial is more likely, given the alternating nature of the sequence: $p_{t^{10000}}(s_{10001}) = p_{t^{10000}}(s_{10001}/t^{10000}) > 0.5 = \bar{\mu}(s_{10001}/t^{10000})$. Accordingly, she would have a strict preference for betting on success in the next trial ex post. Her conditional probabilities would thus change on learning, and she would violate Dynamic Consistency.

Such a decision maker can be straightforwardly modelled by the update rule and framework developed here. Consider the confidence ranking on \mathcal{S}

$$\Xi = \left\{ \{\bar{\mu}\}, \left\{ \pi_{\lambda}^{Markov} : \lambda \in [0, 1] \right\} \cup \{\bar{\mu}\} \right\}$$

where π_{λ}^{Markov} is (the Markov hypothesis) defined by $\pi_{\lambda}^{Markov}(s_{n+1}/s_n) = \pi_{\lambda}^{Markov}(f_{n+1}/f_n) = \lambda$ for all n . This confidence ranking reflects the fact that the parameter space Θ and prior μ over it captures what the decision maker thinks about the sequence of trials she is about to observe: they (or rather the generated predictive) naturally characterise the centre of her confidence ranking. She may be fairly confident in this judgement, and hence use this prior at medium stakes levels (or loss values, for a statistician). However, according to Ξ she is not maximally confident that the process is IID, so there will be confidence levels at which she does not hold this belief. At such levels, the corresponding set of priors contains probability measures that do not correspond to IID processes but, for instance, to Markov processes. For an appropriate $\rho_{t^{10000}}$,

setting reasonable probability thresholds,²⁵ the PT-confidence update of Ξ by t^{10000} is such that $\pi_0^{Markov}(\bullet/t^{10000}) \in \min \Xi_{t^{10000}}$: the decision maker will not have a posterior precise probability of 0.5 for success on the next trial. Indeed, the minimum probability of success over $\min \Xi_{t^{10000}}$ will be greater than 0.5, and will generally depend on the probability threshold set by $\rho_{t^{10000}}$. On this analysis, the confidence approach seems to agree with pre-theoretical intuition. On the one hand, the decision maker sticks with her best-estimate (Bayesian) belief as long as the observations are not very surprising: in the absence of this peculiar pattern, PT-confidence update typically recommends applying Bayesian conditionalisation on $\bar{\mu}$. On the other hand, in the presence of a surprising event (or pattern), she retains only those beliefs held with higher confidence, and moves to the most reasonable conjecture according to those beliefs: in this case, that the process is not IID.

This example can be thought of as a parable of (some) scientific discovery. Prominent discoveries—Fleming’s discovery of penicillin, for instance²⁶—involve noting surprising patterns where one was not expecting them. One certainly would not like to qualify such cases as irrational, and it can be taken as an advantage of our approach that it can capture them comfortably, as the example illustrates. Indeed, it can account for such updates whilst retaining ex ante preferences that are consistent with the initial assumption of an IID process at medium stakes levels; preferences only diverge at high stakes, where lots of confidence is required. By contrast, any Bayesian approach capable of accounting for these sorts of belief change would require the decision maker to place a small probability ex ante on the process not being IID, so as to guarantee that the conditional beliefs remain invariant on learning (Section 4.2.1). This would complicate calculations; behaviourally, it would lead to ex ante preferences which contravene the basic assumption of an IID process. Thus the Bayesian framework cannot accommodate *both* ex ante preferences at medium stakes that are fully compatible with the IID assumption *and* updates that deviate from this assumption in surprising cases. Indeed, a long tradition of experimental and empirical evidence suggests that people do not employ Bayesian updating, especially in the face of ex ante surprising (i.e. low probability) events.²⁷

Similar points hold for crises, such as the financial crisis involved in the opening citation. Evidence suggests that, whilst financial professionals follow Bayesian update in non-crisis periods, they deviate from it in times of crisis (Giacomini et al., 2020). As the example in the Introduction illustrates, confidence update can comfortably capture the opening up to alternative models in the wake of paradigm-shattering events. Moreover, it can also accommodate Bayesian update and decision making in normal times. On the one hand, when the learnt events are not surprising, the investor in the opening example hangs on to his low confidence beliefs, updating them in a Bayesian fashion.²⁸ On the other hand, for medium stakes decisions, his preferences are

²⁵ For instance, take $\rho_{t^{10000}}$ with $\rho_{t^{10000}}(\{\bar{\mu}\}) = 0.05$, $\rho_{t^{10000}}(\{\pi_\lambda^{Markov} : \lambda \in [0, 1]\} \cup \{\bar{\mu}\}) = 0.01$. In this case the conditionalisation of the IID prior $\bar{\mu}$ is no longer held at the bottom of the updated confidence ranking: $\bar{\mu}(\bullet/t^{10000}) \notin \min \Xi_{t^{10000}} = \{\pi_\lambda^{Markov}(\bullet/t^{10000}) : \lambda \in [0, \sqrt[9999]{0.01}]\}$.

²⁶ Fleming noticed a petri dish containing Staphylococci bacteria that had been mistakenly left open was contaminated by blue-green mould from an open window, and that, surprisingly, there was a halo of inhibited bacterial growth around the mould.

²⁷ See for example Kahneman et al. (1982); Grether (1980, 1992); Griffin and Tversky (1992); Camerer et al. (2011) and De Bondt and Thaler (1985, 1987); Gallagher (2014) for experimental and empirical evidence respectively. Note that the points made here hold for both first- and second-order Bayesian approaches.

²⁸ Visually, for non-surprising events, the black and red sets at the ‘Low Confidence’ level in Fig. 1 intersect.




Fig. 2. A Game. (The first number in each pair is Ann’s payoff, the second is Bob’s.)

fully determined by these beliefs—and notably his adopted financial model—although he might entertain a wider range of possibilities at higher confidence levels.

4.2.3. Reasoning in games

Another situation where conditional beliefs may change on update involves surprising or null events in games.

Example 3. Consider the game in Fig. 2 and suppose that Bob thinks that Ann will adopt the Backwards Induction strategy, though admits a small probability ϵ of her making a ‘trembling hand’ mistake at each node. He acts as a Bayesian, and thus places probability $1 - \epsilon$ on her going Out at every node.²⁹ Suppose now that Ann plays In at the first node, then Bob plays In, and then Ann plays In again. By standard Bayesian conditionalisation on the (unexpected, but non-null) event of Ann going in twice, Bob continues to believe that Ann will play the Backwards Induction strategy, and hence go Out at node Ann³. However, given the very small probability of her making two successive mistakes, he might come to reconsider his assumption that she is trying to play the Backwards Induction strategy. He might wonder whether the deviations from Backwards Induction play are intentional: Ann could be aiming for the gain she would get if Bob went In at every node. In other words, he might switch to Forward Induction-style reasoning (Pearce, 1984; Reny, 1992; Stalnaker, 1998; Battigalli and Siniscalchi, 2002). Under this assumption, he would expect Ann to move In at node Ann³. That is, his belief about what Ann would do at node Ann³ after having observed her play In twice differs from his prior conditional belief about what Ann will do if she gets to node Ann³. This is another case where conditional beliefs change on learning (and hence, where Dynamic Consistency is violated).

Given the obvious analogy to Example 2—the assumption that Ann is playing the Backwards Induction strategy plays the role of the IID assumption; the strategy in which she is aiming for both playing In at all nodes plays the role of the alternative Markov hypotheses—it should be no surprise that the confidence approach can comfortably capture the reasoning in this example. Bob’s initial beliefs can be characterised by the centred confidence ranking

$$\Xi = \{ \{ \mu_{BI,\epsilon} \}, \{ \mu_{BI,\epsilon}, \mu_{FI,\epsilon} \} \}$$

where $\mu_{BI,\epsilon}$ is the probability measure over Ann’s play corresponding to the Backwards Induction assumption with ‘trembling hand’ errors— $\mu_{BI,\epsilon}(Out) = 1 - \epsilon > 0.5$ at every node—and $\mu_{FI,\epsilon}$ is the probability measure corresponding to the thesis that Ann is aiming for Bob going In at every node, with ‘trembling hand’ errors— $\mu_{FI,\epsilon}(In) = 1 - \epsilon > 0.5$ at every node. This repre-

²⁹ It is simple to check that Ann’s Backward Induction strategy is to play Out at each node.

sents Bob as a Bayesian with confidence: at the centre of the confidence ranking is the Bayesian belief $\mu_{BI,\epsilon}$, capturing the fact that at intermediate confidence levels, he acts and reasons as a Bayesian accepting that Ann will play the Backward Induction strategy with errors. However, at high levels of confidence, he is not sure of this prediction, entertaining alternative conjectures, and in particular the possibility that Ann is ‘aiming’ for everyone playing In at every node. For appropriate ρ_E , corresponding to appropriate probability thresholds about (the reasoning behind) Ann’s play, confidence update will shift to $\mu_{FI,\epsilon}$ if she makes ‘too many’ mistakes. For instance, if $\rho_{In}(\{\mu_{BI,\epsilon}\}), \rho_{In,In}(\{\mu_{BI,\epsilon}\}), \rho_{In}(\{\mu_{BI,\epsilon}, \mu_{FI,\epsilon}\}), \rho_{In,In}(\{\mu_{BI,\epsilon}, \mu_{FI,\epsilon}\}) \in (\frac{\epsilon}{2}, \epsilon)$,³⁰ then $\min \Xi_{In} = \{\mu_{BI,\epsilon}\}$,³¹ whereas $\min \Xi_{In,In} = \{\mu_{FI,\epsilon}\}$. If Ann goes In once, this can be seen as a mistake, so Bob reasons as a Bayesian and, updating by conditionalisation, sticks with his Backwards Induction assumption. However, if she goes In again, then this is too surprising, and Bob looks to the most reasonable alternative conjecture that he admits as possible, which interprets Ann’s play as intentional. As in the previous example, confidence update can comfortably capture such learning patterns.

Under the confidence analysis, Bob’s update (and reasoning) varies as one would expect with ϵ . For a fixed probability-threshold correspondence ρ_E , as ϵ increases, there may be a value such that, after Ann plays In twice, Bob retains his Backwards Induction assumption: $\min \Xi_{In,In} = \{\mu_{BI,\epsilon}\}$.³² This is as to be expected: if the probability of error is high enough, he need not consider two successive plays of In to be sufficiently surprising, and hence has less reason to doubt his initial beliefs about her strategy. On the other hand, as ϵ decreases with ρ_E fixed, there will be a value below which Bob will interpret Ann’s play as intentional after she plays In just once: $\min \Xi_{In} = \{\mu_{FI,\epsilon}\}$.³³ If he considers a ‘trembling hand’ mistake to be sufficiently unlikely, then seeing just one deviation from the expected Backwards Induction play will be enough to trigger alternative reasoning. This is how Bob would update under this specification in the limit case of no ‘trembling hand’ errors ($\epsilon = 0$). There is thus a continuity in reasoning between very small and zero ex ante probabilities of ‘trembling hand’ errors—that is, between update on surprising and null events.

By contrast, standard approaches retain Bayesian conditionalisation whenever the observed event is non-null: so in the example, whenever $\epsilon > 0$, Bob will hold onto his Backwards Induction assumption no matter how many times Ann plays In. Bayes rule need only be supplemented for cases of update on null events—when $\epsilon = 0$ —and generalisations of probabilities (and Bayesian update) such as conditional probability systems (CPS) or lexicographic conditional probability systems (LCPS) have been proposed for such cases (Rényi, 1955; Myerson, 1986; Blume et al., 1991a,b; Dekel and Siniscalchi, 2015). Since they coincide with Bayesian conditionalisation on non-null events, there is a discontinuity at $\epsilon = 0$: although under the smallest positive probability of error, Bob continues to hold onto the assumption of future Backwards Induction play after several deviations, as soon as the probability hits zero he can change his assumption on update after a single deviation. The continuity supported by the confidence-based approach may seem a

³⁰ In is the event that Ann plays In at the first node; In, In is the event that she plays In at the first two nodes, and so on.

³¹ Note that $\mu_{BI,\epsilon} = \mu_{BI,\epsilon}(\bullet/In)$ and similarly for $\mu_{FI,\epsilon}$ as defined.

³² This occurs whenever $\epsilon^2 \geq \rho_{In}(\{\mu_{BI,\epsilon}\}), \rho_{In,In}(\{\mu_{BI,\epsilon}\})$.

³³ This occurs whenever $\epsilon < \rho_{In}(\{\mu_{BI,\epsilon}\})$.

more desirable property of reasoning in games.³⁴ Whether or not this is so, the example indicates that the confidence framework can cope with update on null events: indeed, the aforementioned generalisations of Bayesian conditionalisation to null events can be recovered as special cases of confidence update, as we now illustrate on CPS's.³⁵

For simplicity, let us assume that the state space \mathcal{S} is finite (and retain all other terminology). A conditional probability system on \mathcal{S} is a map $p^{CPS} : \Sigma \times (\Sigma/\emptyset) \rightarrow [0, 1]$ such that $p^{CPS}(\bullet/E) \in \Delta(\Sigma)$, $p^{CPS}(E/E) = 1$, and $p^{CPS}(E/G) = p^{CPS}(E/F) \cdot p^{CPS}(F/G)$ for all $E, F, G \in \Sigma$ with $E \subseteq F \subseteq G$ and $F \neq \emptyset$. $p^{CPS}(E/S)$ can be thought of as representing prior beliefs. If $p^{CPS}(E/S) > 0$, then $p^{CPS}(\bullet/E)$ is the standard Bayesian conditionalisation of p^{CPS} ; however, $p^{CPS}(\bullet/E)$ is well-defined and non-trivial even when $p^{CPS}(E/S) = 0$. Recall (Section 4.2.1) that a confidence ranking Ξ is centred if it contains a singleton set; in this case, we use p_{Ξ} to denote the member of the singleton set, and call it the *centre* of Ξ .

Proposition 4. *Let p^{CPS} be a conditional probability system on a finite space \mathcal{S} . Then there exists a centred confidence ranking Ξ and a family of functions $(\rho_E)_{E \in \Sigma}$, $\rho_E : \Xi \rightarrow [0, 1]$ such that: i. the centre of Ξ , $p_{\Xi} = p^{CPS}(\bullet/S)$; and ii. for each non-empty event E , Ξ_E , the confidence update of Ξ by E represented by ρ_E is a centred confidence ranking whose centre, p_{Ξ_E} satisfies $p_{\Xi_E}(F) = p^{CPS}(F/E)$ for all $F \in \Sigma$.*

So any decision maker that can be modelled using a CPS can alternatively be modelled using confidence update. Focusing on decisions where the stakes are limited, the decision maker's ex ante and ex post preferences would be precisely as according to the CPS model: in particular confidence update picks out his ex post beliefs properly, even for update on events that are null according to the centre of his confidence ranking. By contrast, his lack of full confidence about his best-guess probability measure (and his relative degree of confidence in the alternatives) comes out in his ex ante preferences under the confidence approach—though not under the CPS approach—in decisions with high or extremely high stakes. On such decisions, his preferences may be non-Bayesian.

This suggests that confidence update, in addition to dealing with update under ambiguity, can comfortably and fruitfully deal with issues arising from update on surprising or null events. Indeed, unlike standard approaches, it offers a uniform treatment of both sorts of update.

5. Discussion

We now briefly consider relationships with other learning paradigms, as well as potential extensions.

Classical and Bayesian statistical reasoning Confidence update subsumes elements of both Classical and Bayesian statistical reasoning. The way it deals with confidence, and in particular the use of probability thresholds over the ex ante probability (or likelihood) of the learnt event under different probability measures, is classical in spirit. The recognition that on learning an

³⁴ We hasten to add that this discussion concerns the reasoning (and update) of one player in a game; evaluating potential consequences for equilibria would require further concepts (e.g. Dekel and Siniscalchi, 2015), and goes beyond the scope of this paper.

³⁵ See for instance Hammond (1994) on the relation with LCPS.

event, one ultimately has to use (some) probabilities conditional on that event is Bayesian. This can be illustrated on Example 1 (Section 4.1).

On the one hand, the penultimate case in the example (involving Dirac measures) reveals a strong analogy to the reasoning in classical statistics: there is a set of parameters (the ex ante set of Dirac measures), and on observation, one can rule out those according to which the observation was too unlikely. The probability threshold in the confidence approach plays a role analogous to the significance level in classical hypothesis testing. However, the confidence approach does not demand a single, fixed significance level. Rather, the update encompasses all relevant significance levels. The level to be used in an ensuing decision is determined by its importance and the decision maker's attitude to choosing on the basis of limited confidence, as represented by his cautiousness coefficient (Section 2.2). In other words, the approach sheds light on how the appropriate significance level should be fixed, revealing the value judgement or taste it corresponds to.

On the other hand, since initial beliefs representable by a Bayesian probability generate a special type of confidence ranking, the confidence update rule can be applied, yielding as posterior beliefs the conditional probability measure (or, more precisely, the confidence ranking whose only element is the singleton containing it). So confidence update coincides with standard Bayesian statistical practice whenever the decision maker holds single-prior beliefs with maximal confidence.

Belief revision Confidence update is also reminiscent of a substantial literature in Artificial Intelligence, logic and philosophy on 'belief revision' (e.g. Gardenfors, 1988), which focuses on belief change in cases where incoming information contradicts initial beliefs. In such cases, there is usually a choice of which among several ex ante beliefs to give up. A popular approach employs the notion of the 'entrenchment' of a belief, and is guided by a maxim similar to ours: hold on to the beliefs that are more 'entrenched', relinquishing those that are less 'entrenched'. This affinity is doubtless related to some of the points made in the preceding sections; indeed, the relevance of belief revision for scientific theory change (Alchourron et al., 1985) and reasoning in games (Stalnaker, 1998) has long been recognised.

However, given the focus on categorical rather than probabilistic beliefs in that literature, it contains, to the best of our knowledge, no rule corresponding to the one proposed here. Moreover, and crucially, they typically do not consider decision. As such, one could consider this paper as developing a decision-theoretic approach to learning that was lacking from the belief revision literature.

Choice and learning An important characteristic of the Bayesian paradigm is the connection between ex ante preferences and update: under it, ex ante and ex post conditional preferences coincide (Section 4.2.1). The current proposal involves a strong, albeit different connection, modulated by the double role of confidence in choice (according to (3)) and learning (via (4)): a decision maker's confidence in a belief regulates both how willing he is to choose on the basis of it and how tenaciously he holds onto it in the face of conflicting information.³⁶ This guarantees that ex post preferences are partially determined by ex ante ones (in particular those held at sufficiently high stakes levels).

³⁶ Note that in the Bayesian paradigm, no single concept plays such a double role: the strength of a Bayesian probability in particular is quite distinct from how tenaciously it is retained on update (e.g. Leitgeb, 2017).

This connection is a central plank of our approach. It draws normative support from the aforementioned intuitions. The relationships it implies between *ex ante* and *ex post* preferences enhance testability, hence lending descriptive clout. And it sets our approach apart from others dealing with null or surprising events. For instance, under the CPS model (Section 4.2.3), *ex ante* preferences impose very few constraints on *ex post* preferences after updating on a null event.

Ortoleva (2012) proposes a ‘Hypothesis Testing’ update rule of Bayesian beliefs which is similar in spirit to the CPS and LCPS models, except that it ‘moves to’ another Bayesian probability when the learnt event is surprising enough (i.e. its *ex ante* probability falls below a threshold), rather than when it is null.³⁷ The rule is motivated by classical hypothesis-testing reasoning, of the sort mentioned above. However, unlike the confidence-based approach, *ex ante* preferences in Ortoleva’s model impose virtually no constraints on the *ex post* preferences an agent will adopt on learning surprising information. In fact, given some underlying technical similarities,³⁸ it may be possible to retrieve the ‘Hypothesis Testing’ rule as a special case of confidence update, via a result similar to Proposition 4 for the CPS model. This may be a way of linking the update to *ex ante* behaviour.

Gilboa et al. (2020) propose a model of choice which combines case-based and expected-utility reasoning, claiming that the former is more appropriate and widespread in the aftermath of surprising events. Since the model is static, it does not draw any link between preferences prior to a (surprising) event and posterior preferences, whereas, as noted, confidence will play a role in relating the two under the approach proposed here.

Extensions and future research Most of the technical assumptions on confidence rankings adopted in Section 3—notably closure and convexity—are inessential to the workings of the update rule. Similar results can be obtained in their absence, albeit with added technicalities to deal, for instance, with the fact that non-convexities do not show up in preferences. Moreover, whilst we have focused on the standard case of update on events, the general logic of the update rule—and in particular the intersection of the sets of probability measures reflecting the information with the *ex ante* confidence ranking—applies for other ‘input formats’, such as information representable by a subset of the probability space (as in Gajdos et al., 2008), a probability assignment for certain events (as in Jeffrey, 1972; Dietrich et al., 2016), or an ambiguous signal (Epstein and Halevy, 2020). Future work could set out the consequences of confidence update in such cases. A final important extension would be to sequential learning situations, as commonly found in statistical decision theory. This would be essential for understanding the long-run implications of the approach, and its consequences in a range of economic applications.

6. Conclusion

This paper proposes a novel update rule under ambiguity. Starting from the intuition that one’s confidence in beliefs has a central role to play in learning, we formulate a model of update of confidence in beliefs, drawing on an existing model of confidence and decision (Hill, 2013). It is based on a simple, but reasonable intuition: when updating in the face of information that conflicts with prior beliefs, *retain* as far as possible those conditional beliefs in which you are

³⁷ Ortoleva (2014) extends the approach to multiple prior beliefs.

³⁸ Specifically: the proof of our Proposition 4 relies on the fact that CPS’s are equivalent to certain orders on the space of probability measures, as are confidence rankings, and Ortoleva’s update is also determined by an order on the probability space (Ortoleva, 2012, Prop 2).

more confident, and relinquish only those in which you have less confidence. A simple and intuitive axiom—Confidence Consistency—characterises a general confidence update rule that conforms to this maxim.

We also characterise a more refined version: probability-threshold confidence update. In a way reminiscent of classical statistical reasoning, it uses a confidence level-dependent threshold to eliminate probability measures that were too ‘wrong’ about the learnt event *ex ante*.

Confidence update can comfortably handle update on complete ignorance, on which standard multiple prior update rules struggle. It provides a general framework that can recover prominent existing update rules as special cases, providing a new perspective on their credentials and relationship. It can also fruitfully deal with update on surprising events, such as crises, and on null events, encompassing the standard game-theoretical tools for the latter as special cases.

Appendix A. Proofs

A.1. Proofs of results in Section 3

Proof of Theorems 1 and 2. We prove Theorem 2. The proof of Theorem 1 is similar. We first show sufficiency of the axioms.

Fix non-null $E \in \Sigma$; since \succeq is non-degenerate by Assumption 1, such events exist. By Assumption 1, there exists a triple (Ξ, D, u) representing \succeq according to (3). For every stakes level σ , let $C^\sigma = D(f)$ for some $f \in \sigma$. It follows from the confidence representation (Hill, 2013) that C^σ represents \succeq^σ (in tandem with u) according to standard maximin EU representation; i.e. \succeq^σ is represented by:

$$V(f) = \min_{p \in C^\sigma} \mathbb{E}_p u(f(s)) \tag{12}$$

As a point of notation, for any $x \in [0, 1]$, we use $[E, x]$ to denote $\{p \in \Delta(\Sigma) : p(E) \geq x\}$.

By Non-degeneracy, \succeq_E is non-degenerate. Moreover, there exists an E -resilient stakes-level σ : if not, by Information-Based Learning, $f \not\sim_E^{\sigma'} c$ for every $f \in \mathcal{A}$, $c \in \mathcal{A}^c$ and stakes level σ' , contradicting the monotonicity of the confidence representation (3).

Lemma 1. *For any E -resilient stakes level σ , there exists $x_E^\sigma \in [0, 1]$ such that \succeq_E^σ is represented by:*

$$V_E^\sigma(f) = \min_{p \in (C^\sigma \cap [E, x_E^\sigma])_E} \mathbb{E}_p u(f) \tag{13}$$

where $(C^\sigma \cap [E, x_E^\sigma])_E$ is as defined in (2). Moreover:

1. if there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ such that $f_{EC} \sim_E^\sigma c$ but $f_{EC} \not\sim^\sigma c$, then there is a unique $x_E^\sigma \in [0, 1]$ having this property;
2. if for all $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, whenever $f_{EC} \sim_E^\sigma c$, then $f_{EC} \sim^\sigma c$, and there exists no $e, d \in \mathcal{A}^c$ with $e > d > c$ and $f_{EE} \geq^\sigma d$, then every $x_E^\sigma \in [0, 1]$ has this property;
3. if for all $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, whenever $f_{EC} \sim_E^\sigma c$, then $f_{EC} \sim^\sigma c$, and for some such $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, there exists $e, d \in \mathcal{A}^c$ with $e > d > c$ and $f_{EE} \geq^\sigma d$, then there exists $\overline{x}_E^\sigma \in [0, 1]$ such that every $x_E^\sigma \in [0, \overline{x}_E^\sigma]$ has this property.

Proof. Fix an E -resilient stakes level σ . For every $f \in \mathcal{A}$, by the representation (Assumption 1), there exists a unique $c \in \mathcal{A}^c$, up to indifference, such that $f_{EC} \sim_E^\sigma c$; consider any

such f and c . For any $e, d \in \mathcal{A}^c$ with $e \succ d \succeq c$ and $f_E e \succeq^\sigma d$, let $\lambda_{e,d;f}$ be the (unique) $\lambda \in [0, 1]$ such that $f_E e \sim^\sigma e_{1-\lambda}d$. (By the E -resilience of σ , $f_E c \preceq^\sigma c$, whence, by the representation, $f_E e \preceq^\sigma e$, so such a λ exists; by the representation, it is unique.) Note that, by definition, for any $p \in \Delta$ such that $\mathbb{E}_p u(f_E e) \geq \mathbb{E}_p u(e_{1-\lambda}d)$ and $\mathbb{E}_p u(e_E d) \geq \mathbb{E}_p u(e_\lambda d)$, we have that $\mathbb{E}_p u((f_E e)_{\frac{1}{2}}(e_E d)) = \mathbb{E}_p u(e_{\frac{1}{2}}(f_E d)) \geq \mathbb{E}_p u((e_{1-\lambda}d)_{\frac{1}{2}}(e_\lambda d)) = \mathbb{E}_p u(e_{\frac{1}{2}}d)$. Let $\Lambda_f^\sigma = \{\lambda_{e,d;f} : e, d \in \mathcal{A}^c, \beta \in (0, 1], f_E e \succeq^\sigma d, e \succ d \succ c\}$, and $\overline{\Lambda}_f^\sigma = \{\lambda_{e,d;f} : e, d \in \mathcal{A}^c, \beta \in (0, 1], f_E e \succeq^\sigma d, e \succ d \succeq c\}$.

Claim 1.

$$\begin{aligned} \Lambda_f^\sigma &= \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } \hat{e} \succ \hat{d}, f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}, \hat{\lambda} > \lambda_{\hat{e},c;f} \right\} \\ &= \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } \hat{e} \succ \hat{d}, f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}, (f_E \hat{e})_{\frac{1}{2}}(\hat{e}_E \hat{d}) \prec_E^\sigma \hat{e}_{\frac{1}{2}} \hat{d} \right\} \end{aligned}$$

and

$$\begin{aligned} \overline{\Lambda}_f^\sigma &= \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } \hat{e} \succ \hat{d}, f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}, \hat{\lambda} \geq \lambda_{\hat{e},c;f} \right\} \\ &= \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } \hat{e} \succ \hat{d}, f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}, (f_E \hat{e})_{\frac{1}{2}}(\hat{e}_E \hat{d}) \leq_E^\sigma \hat{e}_{\frac{1}{2}} \hat{d} \right\}. \end{aligned}$$

Proof. Note firstly that, by the representation, for any $e \succ d, d', \lambda_{e,d;f} > \lambda_{e,d';f}$ if and only if $d \succ d'$. For any $\hat{\lambda} \in [0, 1]$ and $\hat{e}, \hat{d} \in \mathcal{A}^c$ with $\hat{e} \succ \hat{d}$ and $f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}$, if $\hat{\lambda} > \lambda_{\hat{e},c;f}$, then $\hat{d} \succ c$ by the previous observation. So $\Lambda_f^\sigma \supseteq \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e} \succ^\sigma c, \hat{\lambda} > \lambda_{\hat{e},c;f} \right\}$, and similarly for $\overline{\Lambda}_f^\sigma$. Moreover, for such $\hat{\lambda}, \hat{e}, \hat{d}$, it follows from the representation that $(f_E \hat{e})_{\frac{1}{2}}(\hat{e}_E \hat{d}) \prec_E^\sigma \hat{e}_{\frac{1}{2}} \hat{d}$ if and only if $f_E \hat{d} \prec_E^\sigma \hat{d}$, and since $f_E c \sim_E^\sigma c$, this can only be the case if $d \succ c$. So $\Lambda_f^\sigma \supseteq \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } \hat{e} \succ \hat{d}, f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}, (f_E \hat{e})_{\frac{1}{2}}(\hat{e}_E \hat{d}) \prec_E^\sigma \hat{e}_{\frac{1}{2}} \hat{d} \right\}$, and similarly for $\overline{\Lambda}_f^\sigma$. As for the other direction, for any $\hat{e}, \hat{d} \in \mathcal{A}^c$ with $\hat{e} \succ \hat{d} \succ c$ and $f_E \hat{e} \succeq \hat{d}$, if $f_E \hat{e} \sim \hat{e}_{1-\hat{\lambda}} \hat{d}$, then by the previous remark about the ordering of $\lambda_{e,d;f}, \lambda_{e,d';f}, \hat{\lambda} > \lambda_{e,c;f}$; it follows that $\Lambda_f^\sigma \subseteq \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e} \succ^\sigma c, \hat{\lambda} > \lambda_{\hat{e},c;f} \right\}$, and similarly for $\overline{\Lambda}_f^\sigma$. Finally, for any such $\hat{e}, \hat{d} \in \mathcal{A}^c$, by the representation (and in particular C-Independence at a given stakes level) and the fact that $f_E c \sim_E c$, it follows from the $\hat{d} \succ c$ that $f_E \hat{d} \prec_E \hat{d}$, so $(f_E \hat{e})_{\frac{1}{2}}(\hat{e}_E \hat{d}) \prec_E^\sigma \hat{e}_{\frac{1}{2}} \hat{d}$, and hence $\lambda_{\hat{e},\hat{d};f} \in \left\{ \hat{\lambda} \in [0, 1] : \exists \hat{e}, \hat{d} \in \mathcal{A}^c \text{ s.t. } \hat{e} \succ \hat{d}, f_E \hat{e} \sim^\sigma \hat{d}_{\hat{\lambda}} \hat{e}, (f_E \hat{e})_{\frac{1}{2}}(\hat{e}_E \hat{d}) \prec_E^\sigma \hat{e}_{\frac{1}{2}} \hat{d} \right\}$, and similarly for the case of $\hat{d} \succeq c$. This establishes the claim. \square

If, for all $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ such that $f_E c \sim_E^\sigma c, f_E c \sim^\sigma c$, then the result immediately holds with $x_E^\sigma = 0$, so assume henceforth that this is not the case. For clarity, we divide the remainder of the proof into cases.

Case 1. We first consider the case in which there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ with $f_E c \sim_E^\sigma c$ but $f_E c \not\sim^\sigma c$ such that there exists $e, d \in \mathcal{A}^c$ with $e \succ d \succ c$ and $f_E e \succeq^\sigma d$. So Λ_f^σ and $\overline{\Lambda}_f^\sigma$ are non-empty. Since $f_E c \not\sim^\sigma c$, and σ is E -resilient, it follows that $f_E c \prec^\sigma c$; this, in combination with the fact that $f_E c \sim_E^\sigma c$ implies that $f \notin \mathcal{A}^c$. Since, for any $e \in \mathcal{A}^c$

with $f_E e \sim^\sigma e$, $f_E e \succ_E^\sigma e$ by the representation, $0 \notin \Lambda_f^\sigma$. Let $\lambda_f = \inf \Lambda_f^\sigma$. Since, for any $d \succ c$, $\lambda_{\hat{e},d;f} > \lambda_{\hat{e},c;f}$ for all $\hat{e} \in \mathcal{A}^c$, $\lambda_f \notin \Lambda_f^\sigma$, and hence, for every $\bar{c}, \underline{c} \in \mathcal{A}^c$ with $\bar{c} \succ \underline{c}$ and $f_E \bar{c} \sim^\sigma \underline{c}_{\lambda_f} \bar{c}$, it holds that $(f_{\frac{1}{2}} \bar{c})_E (\bar{c}_{\frac{1}{2}} \underline{c}) \succeq_E^\sigma \bar{c}_{\frac{1}{2}} \underline{c}$. Since $\lambda_{e,d;f}$ is continuous in d for every $e \succ c$, for every such e , $\lambda_{e,c;f} \geq \lambda_f$.

We now show that, for every $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$, $\mathbb{E}_p u(f_E c) \geq \mathbb{E}_p u(c)$. First consider any $q \in \mathcal{C}^\sigma \cap \{p \in \Delta(\Sigma) : p(E) > \lambda_f\}$; by the definition of λ_f , there exist $e, d \in \mathcal{A}^c$, with $e \succ d \succ c$, $f_E c \succeq^\sigma d$ and $q(E) \geq \lambda_{e,d;f}$. By the previous remark, since $\mathbb{E}_q u(f_E e) \geq \mathbb{E}_q u(e_{1-\lambda_{e,d;f}} d)$ and $\mathbb{E}_q u(e_E d) \geq \mathbb{E}_q u(e_{\lambda_{e,d;f}} d)$, it follows that $\mathbb{E}_q u((f_E e)_{\frac{1}{2}} (e_E d)) \geq \mathbb{E}_q u(e_{\frac{1}{2}} d)$, and hence, by the linearity of the EU functional, $\mathbb{E}_q u(f_E d) \geq \mathbb{E}_q u(d)$. It follows from the properties of the EU functional that $\mathbb{E}_q u(f_E c) \geq \mathbb{E}_q u(c)$. Since this holds for all $q \in \mathcal{C}^\sigma \cap \{p \in \Delta(\Sigma) : p(E) > \lambda_f\}$, by the continuity of the EU functional, it holds for the closure $\mathcal{C}^\sigma \cap [E, \lambda_f]$, as required.

Now we show that, for each $d \succ c$, there exists $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$ with $\mathbb{E}_p u(f_E d) < u(d)$. For reductio, suppose that there exists $d \succ c$ such that $\mathbb{E}_p u(f_E d) \geq \mathbb{E}_p u(d)$ for all $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$. It follows that $\mathbb{E}_p u(f_E c) > \mathbb{E}_p u(c)$ for all $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$. For each $e \succ c$, consider $I_{e,\lambda_f} = \{p \in \Delta(\Sigma) : \mathbb{E}_p u(f_E e) = \mathbb{E}_p u(e_{1-\lambda_{e,c;f}} c)\} \cap \{p \in \Delta(\Sigma) : p(E) = \lambda_f\}$; since $\mathbb{E}_p u(f_E c) = \mathbb{E}_p u(c)$ for all p in this set (by the previous observation), it follows that $I_{e,\lambda_f} \cap (\mathcal{C}^\sigma \cap [E, \lambda_f]) = \emptyset$ for all such e . Let $\lambda' = \inf \{x \in [0, 1] : \mathbb{E}_p u(f_E c) \geq \mathbb{E}_p u(c), \forall p \in \mathcal{C}^\sigma \cap [E, x]\}$. By the previous observations $\lambda' < \lambda_f$. Moreover, by continuity of the EU functional, there exists $p \in \mathcal{C}^\sigma \cap [E, \lambda']$ such that $\mathbb{E}_p u(f_E c) = \mathbb{E}_p u(c)$. It follows that $I_{e,\lambda'} \cap (\mathcal{C}^\sigma \cap [E, \lambda']) \neq \emptyset$ for at least one $e \succ c$, where $I_{e,\lambda'} = \{p \in \Delta(\Sigma) : \mathbb{E}_p u(f_E e) = \mathbb{E}_p u(e_{1-\lambda_{e,c;f}} c)\} \cap \{p \in \Delta(\Sigma) : p(E) = \lambda'\}$. Since, for any $p \in I_{e,\lambda'}$, $\mathbb{E}_p u((f_E e)_{\frac{1}{2}} (e_E c)) = \mathbb{E}_p u((e_{1-\lambda_{e,c;f}} c)_{\frac{1}{2}} (e_{\lambda'} c)) = u(c_{\frac{1}{2}} (e_{1-(\lambda_{e,c;f}-\lambda')} c))$, and since, for any $p \in I_{e,\lambda'} \cap (\mathcal{C}^\sigma \cap [E, \lambda'])$, $\mathbb{E}_p u(f_E c) \geq \mathbb{E}_p u(c)$, it follows that $\lambda_{e,c;f} = \lambda' < \lambda_f$ for any such e , contradicting the definition of λ_f . So for each $d \succ c$, there exists $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$ with $\mathbb{E}_p u(f_E d) < u(d)$, as required.

Now consider any $f' \in \mathcal{A}$ with $f'_E c' \sim^\sigma c'$. We consider two cases separately.

Case i. First consider the case where $f'_E c' \approx^\sigma c'$. We first treat the case in which there exists $e \in \mathcal{A}^c$ with $e \succ c$ and $f'_E e \succ^\sigma c'$, so, as above, $\Lambda_{f'}^\sigma$ and $\overline{\Lambda_{f'}^\sigma}$ are non-empty. By Probability Consistency, Claim 1 and the previous observations, $\lambda_f \notin \Lambda_{f'}^\sigma$. Applying the same axiom again yields that $\inf \Lambda_{f'}^\sigma \notin \Lambda_{f'}^\sigma$, so $\lambda_f = \inf \Lambda_{f'}^\sigma$. By the arguments used above, $\mathbb{E}_p u(f'_E c') \geq \mathbb{E}_p u(c')$ for all $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$, and, for each $d' \succ c'$, there exists $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$ with $\mathbb{E}_p u(f'_E d') < u(d')$. Now consider the case where, for all $e \in \mathcal{A}^c$, $f'_E e \leq^\sigma c'$. So $\Lambda_{f'}^\sigma = \emptyset$, which by Claim 1, contradicts 3.2.1 and the fact that $\Lambda_f^\sigma \neq \emptyset$, so this case cannot occur.

Case ii. Now consider the case where $f'_E c' \sim^\sigma c'$. So $\mathbb{E}_p u(f'_E c') \geq \mathbb{E}_p u(c')$ for all $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$. If there exists $e, d \in \mathcal{A}^c$ with $e \succ d \succ c'$ and $f'_E e \geq^\sigma d$, then $\Lambda_{f'}^\sigma \neq \emptyset$. By Probability Consistency and the last characterisation of $\Lambda_{f'}^\sigma$ in Claim 1, $\lambda_f < \lambda$ for all $\lambda \in \Lambda_{f'}^\sigma$. By an argument similar to that used above that, for each $d' \succ c'$, there exists $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$ with $\mathbb{E}_p u(f'_E d') < u(d')$. If there exists no $e, d \in \mathcal{A}^c$ with $e \succ d \succ c'$ and $f'_E e \geq^\sigma d$, then $f'_E e \sim f'_E c' \sim^\sigma c'$ for all $e \in \mathcal{A}^c$ with $e \succ c'$. It follows from the representation that there ex-

ists $p \in \mathcal{C}^\sigma$ with $\mathbb{E}_p u(f_E d') = u(c') < u(d')$ for all $d' \succ c'$ and $p(E) = 1$; since $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$, for every $d' \succ c'$, there exists $p \in \mathcal{C}^\sigma \cap [E, \lambda_f]$ with $\mathbb{E}_p u(f_E d') < u(d')$.

Case 2. Now we consider the case in which there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ such that $f_{EC} \sim_E^\sigma c$ but $f_{EC} \not\sim^\sigma c$, and for all such f, c , $f_{Ee} \preceq^\sigma c$ for all $e \in \mathcal{A}^c$. By Null consistency, for each such f, c , there exists $e \in \mathcal{A}^c$ with $f_{Ee} \sim^\sigma c$. Since $f_{Ee'} \sim^\sigma f_{Ee}$ for any $e' \succ e$ and any such f, c , it follows from the representation that there exists $p \in \mathcal{C}^\sigma$ with $\mathbb{E}_p u(f_{Ee}) = u(c)$ and $p(E) = 1$ and that, for any other $q \in \mathcal{C}^\sigma$ with $q(E) = 1$, $\mathbb{E}_q u(f_{Ee}) \geq u(c)$. It thus follows that for all $p \in \mathcal{C}^\sigma \cap [E, 1]$, $\mathbb{E}_p u(f_{EC}) \geq \mathbb{E}_p u(c)$. Moreover, for every $d \succ c$, if $\mathbb{E}_p u(f_{Ed}) \geq \mathbb{E}_p u(d)$ for all $p \in \mathcal{C}^\sigma \cap [E, 1]$, then $f_{Ed} \succ^\sigma c$, contradicting the definition of the case; so for each $d \succ c$, there exists $p \in \mathcal{C}^\sigma \cap [E, 1]$ with $\mathbb{E}_p u(f_{Ed}) < u(d)$. Now consider any $f' \in \mathcal{A}$ with $f'_E c' \sim_E^\sigma c'$ and $f'_E c' \not\sim^\sigma c'$. If there exists $e, d \in \mathcal{A}^c$ with $e \succ d \succ c'$ and $f'_E e \succeq^\sigma d$, then $\Lambda_{f'}^\sigma \cap [0, 1] \neq \emptyset$. By Probability Consistency and the last characterisation of Λ_f in Claim 1, it follows that, for every $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ such that $f_{EC} \sim_E^\sigma c$ but $f_{EC} \not\sim^\sigma c$, and $f_{Ee} \preceq^\sigma c$ for all $e \in \mathcal{A}^c$, there exists $e' \succ d' \succ c$ with $f_{Ee'} \succeq^\sigma d' \succ c$, which is a contradiction. So for every $f' \in \mathcal{A}$ with $f'_E c' \sim_E^\sigma c'$ and $f'_E c' \not\sim^\sigma c'$, $f'_E e \sim^\sigma c'$ for all $e \in \mathcal{A}^c$ with $e \succ c'$. It follows from the representation that there exists $p \in \mathcal{C}^\sigma$ with $\mathbb{E}_p u(f_E d') = u(c') < u(d')$ for all $d' \succ c'$ and $p(E) = 1$; since $p \in \mathcal{C}^\sigma \cap [E, 1]$, for every $d' \succ c'$, there exists $p \in \mathcal{C}^\sigma \cap [E, 1]$ with $\mathbb{E}_p u(f_E d') < u(d')$.

Let $x_E^\sigma = \lambda_f$ in Case 1 and $x_E^\sigma = 1$ in Case 2. By the previous observations, for every $f \in \mathcal{A}$, $\min_{p \in \mathcal{C}^\sigma \cap [E, x_E^\sigma]} \mathbb{E}_p u(f_{EC}) \geq u(c)$, where $f_{EC} \sim_E^\sigma c$, and for any $d \succ c$, $\min_{p \in \mathcal{C}^\sigma \cap [E, x_E^\sigma]} \mathbb{E}_p u(f_E d) < u(d)$. It follows from the continuity of the maximin-EU functional that $\min_{p \in \mathcal{C}^\sigma \cap [E, x_E^\sigma]} \mathbb{E}_p u(f_{EC}) = u(c)$ for all $f \in \mathcal{A}$ with $f_{EC} \sim_E^\sigma c$. By Consequentialism, for every $f \in \mathcal{A}$, $f \sim_E^\sigma c$ for $c \in \mathcal{A}^c$ such that $f_{EC} \sim_E^\sigma c$, so the preferences \succeq_E^σ are represented by $V(f) = u(c)$ such that $f_{EC} \sim_E^\sigma c$. Since:

$$\begin{aligned} \min_{p \in \mathcal{C}^\sigma \cap [E, x_E^\sigma]} \mathbb{E}_p u(f_{EC}) = u(c) &\Leftrightarrow \min_{p \in \mathcal{C}^\sigma \cap [E, x_E^\sigma]} (p(E)(\mathbb{E}_{p(\bullet/E)} u(f)) + (1 - p(E))u(c)) \\ &= u(c) \\ &\Leftrightarrow \min_{p \in \mathcal{C}^\sigma \cap [E, x_E^\sigma]} \mathbb{E}_{p(\bullet/E)} u(f) = u(c) \\ &\Leftrightarrow \min_{p \in (\mathcal{C}^\sigma \cap [E, x_E^\sigma])_E} \mathbb{E}_p u(f) = u(c) \end{aligned}$$

This establishes the representation.

As concerns the uniqueness of x_E^σ , it is clear from the proof that, if there exist $f \in \mathcal{A}$ with $c \in \mathcal{A}^c$ such that $f_{EC} \sim_E^\sigma c$ but $f_{EC} \not\sim^\sigma c$, then $x_E^\sigma = \inf \Lambda_f^\sigma$ for any such $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ in Case 1, and $x_E^\sigma = 1$ if Case 2 holds. Since Λ_f^σ is uniquely defined on the basis of preferences, this implies that x_E^σ is unique. If $f_{EC} \sim_E^\sigma c$ whenever $f_{EC} \sim^\sigma c$, and for no such f, c there exists $e, d \in \mathcal{A}^c$ with $e \succ d \succ c$ and $f_{Ee} \succeq^\sigma d$, then by the analysis of this case 1.ii., $\min_{p \in \mathcal{C}^\sigma \cap [E, x]} \mathbb{E}_p u(f_{EC}) = u(c)$ iff $f_{EC} \sim_E^\sigma c$, for all $x \in [0, 1]$, as required. Finally, if $f_{EC} \sim_E^\sigma c$ whenever $f_{EC} \sim^\sigma c$ but for some $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, there exists $e, d \in \mathcal{A}^c$ with $e \succ d \succ c$ and $f_{Ee} \succeq^\sigma d$, then by the observations about Λ_f^σ (case 1.ii.), $\min_{p \in \mathcal{C}^\sigma \cap [E, x]} \mathbb{E}_p u(g_{Ed}) = u(d)$ iff $g_{Ed} \sim_E^\sigma d$, whenever $x < \lambda$ for all $\lambda \in \Lambda_f^\sigma$, as required. \square

Define the function ϕ_E relating E -resilient stakes levels to values in $[0, 1]$ as follows:

1. If σ satisfies the conditions of clause 1. of Lemma 1, then $\phi_E(\sigma) = x_E^\sigma$ such that (13) holds.
2. If σ satisfies the conditions of clause 2. of Lemma 1, then $\phi_E(\sigma) = \sup\{\phi_E(\sigma') : \sigma' > \sigma\}$.³⁹
3. If σ satisfies the conditions of clause 3. of Lemma 1, then $\phi_E(\sigma) = \max\left\{\sup\{\phi_E(\sigma') : \sigma' > \sigma\}, \overline{x_E^\sigma}\right\}$, where $\overline{x_E^\sigma}$ is as in Lemma 1.

By definition and Lemma 1, (13) holds for $\phi_E(\sigma)$ for every E -resilient stakes level σ .

Claim 2. For every E -resilient σ', σ'' with $\sigma'' > \sigma'$, $\phi_E(\sigma'') \leq \phi_E(\sigma')$.

Proof. Let $\Lambda_{f'}^{\sigma'}$ and $\Lambda_{f''}^{\sigma''}$ be defined as in the proof of Lemma 1, for appropriate f', f'' . By the proof of that Lemma, if the stakes levels σ', σ'' satisfy the conditions of clause 1. (i.e. there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ with $f_{EC} \sim^{\sigma'} c$ but $f_{EC} \not\sim^{\sigma''} c$ and similarly for σ''), then $x_E^{\sigma'} = \inf \Lambda_{f'}^{\sigma'}$ (under case 1 in the proof of the Lemma) or $x_E^{\sigma'} = 1$ (in case 2), and similarly for $x_E^{\sigma''}$. By Probability Consistency and Claim 1, for any $\lambda \notin \Lambda_{f''}^{\sigma''}, \lambda \notin \Lambda_{f'}^{\sigma'}$, so if $x_E^{\sigma''} = 1$, then $x_E^{\sigma'} = 1$ (both stakes levels are in case 2), and if $x_E^{\sigma''} = \inf \Lambda_{f''}^{\sigma''} < 1$, $x_E^{\sigma'} = \min\left\{\inf \Lambda_{f'}^{\sigma'}, 1\right\} \geq \inf \Lambda_{f''}^{\sigma''} = x_E^{\sigma''}$. If σ' satisfies the conditions of clause 1 and is in case 1 of Lemma 1 (so $x_E^{\sigma'} = \inf \Lambda_{f'}^{\sigma'} < 1$) and σ'' satisfies the conditions of clause 3. (in particular, for some $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$, there exists $e, d \in \mathcal{A}^c$ with $e > d > c$ and $f_{Ee} \succeq^{\sigma''} d$), $\overline{x_E^{\sigma''}} = \inf \Lambda_{f''}^{\sigma''}$ for appropriate f'' , and $x_E^{\sigma'} = \inf \Lambda_{f'}^{\sigma'} \geq \inf \Lambda_{f''}^{\sigma''} = \overline{x_E^{\sigma''}}$ by Probability Consistency and Claim 1, which implies, in the light of the previous analysis of case of clause 1, that $\phi_E(\sigma'') \leq \phi_E(\sigma')$. If σ' satisfies the conditions of clause 1 and is in case 2 of Lemma 1 (so $x_E^{\sigma'} = 1$), then by Probability Consistency and Claim 1 and the argument in case 2 of Lemma 1, σ'' does not satisfy the conditions of clause 3. Given the previous two cases, if σ' satisfies the conditions of clause 1. and σ'' satisfies the conditions of clause 2., then it follows from clause 2. and the fact that $\phi_E(\sigma_1) \leq \phi_E(\sigma_2)$ for all $\sigma_1 > \sigma_2$ satisfying the conditions of clause 1 or 3, that $\phi_E(\sigma'') \leq \phi_E(\sigma')$. If σ'' satisfies the conditions of clauses 2 or 3, then the result is immediate. \square

For every E -resilient σ , let $x_E^\sigma = \phi_E(\sigma)$. Let $\mathcal{D} = \bigcap_{\sigma' \text{ } E\text{-resilient}} \overline{(\mathcal{C}^{\sigma'} \cap [E, x_E^{\sigma'}])_E}$ and $y_E = \sup_{\sigma \text{ } E\text{-resilient}} \phi_E(\sigma)$. As noted above, there exists an E -resilient stakes level, so $\mathcal{D} \neq \emptyset$. By Confidence Consistency, for any stakes level σ'' that is not E -resilient, $\sigma' > \sigma''$ for every E -resilient stakes level σ' . It follows from the confidence representation (3) that $\mathcal{C}'' \subseteq \mathcal{D}$ for every \mathcal{C}'' representing $\succeq_E^{\sigma''}$ according to (12).

Claim 3. Under Information-based Learning, for any stakes level σ'' , if σ'' is not E -resilient, then $\succeq_E^{\sigma''}$ is represented by \mathcal{D} .

Proof. Let σ'' be a stakes level that is not E -resilient, and let \mathcal{C}'' be a closed convex set representing $\succeq_E^{\sigma''}$ according to (12). (Such a set exists by the representation (3).) As noted above,

³⁹ We adopt the convention that the infimum over the empty set is 1.

$C'' \subseteq \mathcal{D}$). Suppose that the inverse containment does not hold, so there exists $p \in \text{convcl}(\mathcal{D}) \setminus C''$. By a separating hyperplane argument, there exists $f \in \mathcal{A}$, $c \in \mathcal{A}^c$ such that $\mathbb{E}_q u(f) \geq u(c)$ for all $q \in C''$ whereas $\mathbb{E}_p u(f) < u(c)$. It follows that $f \not\geq_E^{\sigma'} c$ for all E -resilient σ' but $f \geq_E^{\sigma''} c$, contradicting Information-based Learning. So $C'' = \mathcal{D}$ and \mathcal{D} represents $\geq_E^{\sigma''}$, as required. \square

Claim 4. For any stakes level σ'' that is not E -resilient and any C'' representing $\geq_E^{\sigma''}$ according to (12), $C'' \cap \{p \in \Delta(\Sigma) : p(E) \geq y_E\} = \emptyset$.

Proof. Consider a non- E -resilient σ'' , and let C'' represent $\geq_E^{\sigma''}$. By Claim 3, $\mathcal{D} = \overline{\bigcap_{\sigma' \text{ } E\text{-resilient}} (C^{\sigma'} \cap [E, x_E^{\sigma'}])_E} = \overline{(\bigcap_{\sigma' \text{ } E\text{-resilient}} (C^{\sigma'} \cap [E, x_E^{\sigma'}]))_E}$ represents $\geq_E^{\sigma''}$; however, by Confidence Consistency and the confidence representation, $C'' \subseteq \bigcap_{\sigma' \text{ } E\text{-resilient}} C^{\sigma'}$. So if $C'' \cap \{p \in \Delta(\Sigma) : p(E) \geq y_E\} \neq \emptyset$, then $C'' \cap \bigcap_{\sigma' \text{ } E\text{-resilient}} (C^{\sigma'} \cap [E, x_E^{\sigma'}]) = C'' \cap \bigcap_{\sigma' \text{ } E\text{-resilient}} C^{\sigma'} \cap \bigcap_{\sigma' \text{ } E\text{-resilient}} [E, x_E^{\sigma'}] = C'' \cap \bigcap_{\sigma' \text{ } E\text{-resilient}} C^{\sigma'} \cap [E, y_E] \neq \emptyset$, and hence, for every $f \in \mathcal{A}$, $c \in \mathcal{A}^c$, if $f_{EC} \geq_E^{\sigma''} c$, then $f_{EC} \geq_E^{\sigma'} c$ by the reasoning in the proof of Lemma 1, contradicting the assumption that σ'' is not E -resilient. Hence $C'' \cap \{p \in \Delta(\Sigma) : p(E) \geq y_E\} = \emptyset$ as required. \square

Define $\rho_E : \Xi \Rightarrow [0, 1]$ as follows⁴⁰:

$$\rho_E(\mathcal{C}) = \begin{cases} \{x_E^\sigma : D^{-1}(\mathcal{C}) \cap \sigma \neq \emptyset, \sigma \text{ } E\text{-resilient}\} & \text{if } \exists E\text{-resilient } \sigma \\ & \text{s.t. } D^{-1}(\mathcal{C}) \cap \sigma \neq \emptyset \\ y_E^\sigma & \text{otherwise} \end{cases} \quad (14)$$

Since $x_E^{\sigma'} \geq x_E^{\sigma''}$ whenever $\sigma' \leq \sigma''$ with σ', σ'' E -resilient, and since D respects \geq , ρ_E is a decreasing correspondence. By the fact that, for every E -resilient σ , \geq_E^σ is represented by $(C^{\sigma'} \cap [E, x_E^{\sigma'}])_E$ and by Claims 3 and 4, Ξ_E , defined with respect to ρ_E as in (5), represents \geq_E . Hence \geq_E is a confidence update of \geq , as required.

Necessity of the axioms is straightforward, given, in the cases of Probability Consistency and Null consistency, the insights involved in Lemma 1 and its proof. \square

Proof of Proposition 1. Lemma 1 implies that x_E^σ is uniquely defined under clause 1, which immediately implies the second part of the uniqueness clause, taking \mathcal{C} to be the \mathcal{D} in the proof of the Theorem. As for the first part, it follows from the fact that ϕ_E in the proof of Theorem 2 was defined to take the highest admissible value for each stakes level, and the fact that there a unique highest admissible value for each stakes level, by Lemma 1. \square

Proof of Proposition 2. Use the same reasoning as the proof of Theorem 2 and Proposition 1, relying on the following strengthening of Lemma 1.

Lemma 2. Under the conditions in Theorem 2 and Strong Probability Consistency, for all non-null $E, F \in \Sigma$ and for every stakes level σ that is both E - and F -resilient, there exists $x^\sigma \in [0, 1]$ such that \geq_E^σ and \geq_F^σ are represented according to (13) with x^σ . (I.e. $V_E^\sigma(f) =$

⁴⁰ Recall from Section 2.3 that stakes levels are defined as sets (equivalence classes) of acts.

$\min_{p \in (\mathcal{C}^\sigma \cap [E, x^\sigma])_E} \mathbb{E}_p u(f)$ represents \succeq_E^σ and $V_F^\sigma(f) = \min_{p \in (\mathcal{C}^\sigma \cap [F, x^\sigma])_F} \mathbb{E}_p u(f)$ represents \succeq_F^σ .) Moreover, the uniqueness of x^σ is as in Lemma 1.

Proof. The proof employs the same reasoning as the proof of Lemma 1, with the definition of cases by (for instance) “there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ with $f_{EC} \sim_E^\sigma c$ but $f_{EC} \not\sim^\sigma c$ such that there exists $e \in \mathcal{A}^c$ with $e \succ c$ and $f_{EE} \succ^\sigma c$ ” replaced by “there exists $f \in \mathcal{A}$ and $c \in \mathcal{A}^c$ with $f_{EC} \sim_E^\sigma c$, $f_{EC} \not\sim^\sigma c$ and $f_{EE} \succ^\sigma c$ for some $e \in \mathcal{A}^c$ with $e \succ c$, or with $f_{FC} \sim_F^\sigma c$, $f_{FC} \not\sim^\sigma c$ and $f_{FE} \succ^\sigma c$ for some $e \in \mathcal{A}^c$ with $e \succ c$ ” (and similarly for the other cases). \square

\square

Proof of Proposition 3. Fix a non-null event E , and let Ξ_E , respectively Ξ'_E be the confidence rankings and D_E and D'_E the cautiousness coefficients representing \succeq_E and \succeq'_E and obtained by confidence update according to Theorem 2. Let ϕ_E and ϕ'_E be as defined prior to Claim 2 in the proof of Theorem 2 for decision makers \succeq and \succeq' respectively. By (Hill, 2013, Thm 2 and the arguments used in its proof), (i) iff for every stakes level σ that is E -resilient according to \succeq , $D'_E(f) \supseteq D_E(f)$ for every $f \in \sigma$. Note that it follows that any such stakes level is also E -resilient according to \succeq' . Since, by Theorem 2 and its proof, $D_E(f) = \overline{(D(f) \cap [E, \phi_E(\sigma_f)])}_E$, and similarly for $D'_E(f)$, the previous containment holds iff $\phi_E(\sigma) \geq \phi'_E(\sigma)$ for every such stakes level σ . By the definition of the maximal correspondences ρ_E and ρ'_E representing the updates, this holds iff (ii), as required. \square

A.2. Proofs of other results

Proof of Proposition 4. As is well-known (Hammond, 1994), when the state space is finite, p^{CPS} is equivalent to a sequence (p_1, \dots, p_n) of (ordinary) probability measures, with disjoint supports, in the following sense: for every $E_1, E_2 \in \Sigma$ with $E_2 \neq \emptyset$, $p^{CPS}(E_1/E_2) = p_j(E_1/E_2)$ where $p_j(E_2) \neq 0$ and $p_k(E_2) = 0$ for all $k < j$. Define the confidence ranking $\Xi(p^{CPS}) = \{\{p_i : i \leq k\} : k = 1, \dots, n\}$. This is a well-defined min-closed confidence ranking. Taking, for each $E \in \Sigma$, ρ_E with $\rho_E(\mathcal{C}) = 0$ for all $\mathcal{C} \in \Xi(p^{CPS})$, it is clear that, for every $E \in \Sigma$, the confidence update $\Xi(p^{CPS})_E = \{\{p_i(\bullet/E) : p_i(E) > 0\}$, whose centre $p_{\Xi(p^{CPS})_E} = p_j(\bullet/E)$ where $p_k(E) = 0$ for all $k < j$. Hence $p_{\Xi(p^{CPS})_E} = p^{CPS}(\bullet/E)$, and the confidence update exhibits the same conditional probabilities are the conditional probability system p^{CPS} , as required. \square

References

Alchourron, C.E., Gärdenfors, P., Makinson, D., 1985. The logic of theory change: partial meet contraction and revision functions. *J. Symb. Log.* 50, 510–530.
 Anscombe, F.J., Aumann, R.J., 1963. A definition of subjective probability. *Ann. Math. Stat.* 34, 199–205.
 Baillon, A., Bleichrodt, H., Keskin, U., l’Haridon, O., Li, C., 2018. The effect of learning on ambiguity attitudes. *Manag. Sci.* 64, 2181–2198.
 Battigalli, P., Siniscalchi, M., 2002. Strong belief and forward induction reasoning. *J. Econ. Theory* 106, 356–391.
 Bland, J.R., Rosokha, Y., 2019. Learning Under Uncertainty with Multiple Priors: Experimental Investigation. Available at SSRN 3419020.
 Blume, L., Brandenburger, A., Dekel, E., 1991a. Lexicographic probabilities and choice under uncertainty. *Econometrica* 59, 61–79.
 Blume, L., Brandenburger, A., Dekel, E., 1991b. Lexicographic probabilities and equilibrium refinements. *Econometrica*, 81–98.
 Bradley, R., 2005. Radical probabilism and mental kinematics. *Philos. Sci.* 72, 342–364.
 Bradley, R., 2017. *Decision Theory with a Human Face*. Cambridge University Press.

- Bradley, S., 2014. Imprecise probabilities. In: Zalta, E.N. (Ed.), *The Stanford Encyclopedia of Philosophy*, winter 2014 ed.
- Caballero, R.J., Krishnamurthy, A., 2008. Collective risk management in a flight to quality episode. *J. Finance* 63, 2195–2230.
- Camerer, C.F., Loewenstein, G., Rabin, M., 2011. *Advances in Behavioral Economics*. Princeton University Press.
- Cohen, M., Gilboa, I., Jaffray, J.-Y., Schmeidler, D., 2000. An experimental study of updating ambiguous beliefs. *Risk Decis. Policy* 5, 123–133. publisher: Cambridge University Press.
- Cox, L.A.T., 2012. Confronting deep uncertainties in risk analysis. *Risk Anal.* 32, 1607–1629.
- De Bondt, W.F., Thaler, R., 1985. Does the stock market overreact? *J. Finance* 40, 793–805.
- De Bondt, W.F., Thaler, R.H., 1987. Further evidence on investor overreaction and stock market seasonality. *J. Finance*, 557–581.
- Dekel, E., Siniscalchi, M., 2015. Epistemic game theory. In: *Handbook of Game Theory with Economic Applications*, vol. 4. Elsevier, pp. 619–702.
- Dempster, A.P., 1967. Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Stat.* 38, 325–339.
- Dietrich, F., List, C., Bradley, R., 2016. Belief revision generalized: a joint characterization of Bayes' and Jeffrey's rules. *J. Econ. Theory* 162, 352–371.
- Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. *Q. J. Econ.* 75, 643–669.
- Epstein, L.G., Halevy, Y., 2020. Hard-to-interpret signals. Tech. Rep. University of Boston.
- Epstein, L.G., Le Breton, M., 1993. Dynamically consistent beliefs must be Bayesian. *J. Econ. Theory* 61, 1–22.
- Epstein, L.G., Schneider, M., 2003. Recursive multiple-priors. *J. Econ. Theory* 113, 1–31.
- Epstein, L.G., Schneider, M., 2007. Learning under ambiguity. *Rev. Econ. Stud.* 74, 1275–1303.
- Gajdos, T., Hayashi, T., Tallon, J.-M., Vergnaud, J.-C., 2008. Attitude toward imprecise information. *J. Econ. Theory* 140, 27–65.
- Gallagher, J., 2014. Learning about an infrequent event: evidence from flood insurance take-up in the United States. *Am. Econ. J. Appl. Econ.* 6, 206–233.
- Gärdenfors, P., 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, MA.
- Ghirardato, P., 2002. Revisiting Savage in a conditional world. *Econ. Theory* 20, 83–92.
- Ghirardato, P., Marinacci, M., 2002. Ambiguity made precise: a comparative foundation. *J. Econ. Theory* 102, 251–289.
- Giacomini, R., Skreta, V., Turen, J., 2020. Heterogeneity, inattention, and Bayesian updates. *Am. Econ. J. Macroecon.* 12, 282–309.
- Gilboa, I., Marinacci, M., 2013. Ambiguity and the Bayesian paradigm. In: *Advances in Economics and Econometrics: Volume 1, Economic Theory: Tenth World Congress*, vol. 49. Cambridge University Press, p. 179.
- Gilboa, I., Minardi, S., Samuelson, L., 2020. Theories and cases in decisions under uncertainty. *Games Econ. Behav.* 123, 22–40.
- Gilboa, I., Postlewaite, A., Schmeidler, D., 2009. Is it always rational to satisfy Savage's axioms? *Econ. Philos.*, 285–296.
- Gilboa, I., Postlewaite, A., Schmeidler, D., 2012. Rationality of belief. *Synthese* 187, 11–31.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *J. Math. Econ.* 18, 141–153.
- Gilboa, I., Schmeidler, D., 1993. Updating ambiguous beliefs. *J. Econ. Theory* 59, 33–49.
- Grether, D.M., 1980. Bayes rule as a descriptive model: the representativeness heuristic. *Q. J. Econ.* 95, 537–557.
- Grether, D.M., 1992. Testing Bayes rule and the representativeness heuristic: some experimental evidence. *J. Econ. Behav. Organ.* 17, 31–57.
- Griffin, D., Tversky, A., 1992. The weighing of evidence and the determinants of confidence. *Cogn. Psychol.* 24, 411–435.
- Gul, F., Pesendorfer, W., 2021. Evaluating ambiguous random variables from Choquet to maxmin expected utility. *J. Econ. Theory* 192, 105129.
- Hammond, P.J., 1994. Elementary non-Archimedean representations of probability for decision theory and games. In: *Patrick Suppes: Scientific Philosopher*. Springer, pp. 25–61.
- Hanany, E., Klibanoff, P., 2007. Updating preferences with multiple priors. *Theor. Econ.* 2, 261–298.
- Hill, B., 2013. Confidence and decision. *Games Econ. Behav.* 82, 675–692.
- Hill, B., 2015. Confidence as a Source of Deferral. HEC Paris Research Paper ECO/SCD-2014-1060. HEC Paris. April 11, 2015 version.
- Hill, B., 2016. Incomplete preferences and confidence. *J. Math. Econ.* 65, 83–103.
- Hill, B., 2019. Confidence in beliefs and rational decision making. *Econ. Philos.* 35, 223–258.
- Hill, B., 2020. Dynamic consistency and ambiguity: a reappraisal. *Games Econ. Behav.* 120, 289–310.
- Ilut, C.L., Schneider, M., 2014. Ambiguous business cycles. *Am. Econ. Rev.* 104, 2368–2399.
- Jeffrey, R.C., 1972. *The Logic of Decision*, 2nd ed. University of Chicago Press, Chicago.

- Jeffrey, R.C., 1992. *Probability and the Art of Judgment*. Cambridge University Press.
- Joyce, J.M., 2010. A defense of imprecise credences in inference and decision making. *Philos. Perspect.* 24, 281–323.
- Kahneman, D., Slovic, P., Tversky, A., 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Kovach, M., 2015. *Ambiguity and Partial Bayesian Updating*. Mimeo. Caltech.
- Leitgeb, H., 2017. *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.
- Levi, I., 1974. On indeterminate probabilities. *J. Philos.* 71, 391–418.
- Manski, C.F., 2013. *Public Policy in an Uncertain World: Analysis and Decisions*. Harvard University Press, Cambridge, Mass.
- Manski, C.F., Nagin, D.S., 1998. Bounding disagreements about treatment effects: a case study of sentencing and recidivism. *Sociol. Method.* 28, 99–137.
- Myerson, R.B., 1986. Multistage games with communication. *Econometrica*, 323–358.
- Ortoleva, P., 2012. Modeling the change of paradigm: non-Bayesian reactions to unexpected news. *Am. Econ. Rev.* 102, 2410–2436.
- Ortoleva, P., 2014. Hypothesis testing and ambiguity aversion. *Riv. Polit. Econ.* 3, 45–64.
- Pearce, D.G., 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 1029–1050.
- Pires, C.P., 2002. A rule for updating ambiguous beliefs. *Theory Decis.* 53, 137–152.
- Reny, P.J., 1992. Backward induction, normal form perfection and explicable equilibria. *Econometrica*, 627–649.
- Rényi, A., 1955. On a new axiomatic theory of probability. *Acta Math. Hung.* 6, 285–335.
- Stalnaker, R., 1998. Belief revision in games: forward and backward induction. *Math. Soc. Sci.* 36, 31–56.
- Vallinder, A., 2017. Imprecise Bayesianism and global belief inertia. *Br. J. Philos. Sci.* 69, 1205–1230.
- Walley, P., 1991. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.