

REPORT DOCUMENTATION PAGE			2		Form Approved OMB NO. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)		
		New Reprint		-		
4. TITLE AND SUBTITLE Conditional belief types				5a. CONTRACT NUMBER		
				W911NF-14-1-0017		
				5b. GRANT NUMBER		
6. AUTHORS Joseph Y. Halpern, Dov Samet, Alfredo Di Tillio				5c. PROGRAM ELEMENT NUMBER		
				611102		
				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES				8. PERFORMING ORGANIZATION REPORT NUMBER		
Cornell University 373 Pine Tree Road Ithaca, NY 14850 -2820						
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 64452-NS.5		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.						
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.						
14. ABSTRACT We study type spaces where a player's type at a state is a conditional probability on the space. We axiomatize these spaces using conditional belief operators, examining three additional axioms of increasing strength. First, introspection, which requires the agent to be unconditionally certain of her beliefs. Second, echo, according to which the unconditional beliefs implied by the condition must be held given the condition. Third, determination, which says that the conditional beliefs are the unconditional beliefs that are conditionally certain. Echo implies that conditioning on an event is the same as conditioning on the event being certain, which formalizes the standard						
15. SUBJECT TERMS type space, hypothetical knowledge, conditional probability, counterfactuals						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Joseph Halpern	
UU	UU	UU	UU		19b. TELEPHONE NUMBER 607-255-9562	

ABSTRACT

We study type spaces where a player's type at a state is a conditional probability on the space. We axiomatize these spaces using conditional belief operators, examining three additional axioms of increasing strength. First, introspection, which requires the agent to be unconditionally certain of her beliefs. Second, echo, according to which the unconditional beliefs implied by the condition must be held given the condition. Third, determination, which says that the conditional beliefs are the unconditional beliefs that are conditionally certain. Echo implies that conditioning on an event is the same as conditioning on the event being certain, which formalizes the standard informal interpretation of conditional probability. The game-theoretic application of our model, discussed within an example, sheds light on a number of issues in the analysis of extensive form games. Type spaces are closely related to the sphere models of counterfactual conditionals and to models of hypothetical knowledge.

REPORT DOCUMENTATION PAGE (SF298) (Continuation Sheet)

Continuation for Block 13

ARO Report Number 64452.5-NS

Conditional belief types ...

Block 13: Supplementary Note

© 2014 . Published in Games and Economic Behavior, Vol. Ed. 0 87, (0) (2014), (, (0). DoD Components reserve a royalty-free, nonexclusive and irrevocable right to reproduce, publish, or otherwise use the work for Federal purposes, and to authorize others to do so (DODGARS §32.36). The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.

Approved for public release; distribution is unlimited.

Conditional belief types*

Alfredo Di Tillio[†] Joseph Y. Halpern[‡] Dov Samet[§]

June 1, 2014

Abstract

We study type spaces where a player's type at a state is a conditional probability on the space. We axiomatize these spaces using conditional belief operators, examining three additional axioms of increasing strength. First, *introspection*, which requires the agent to be unconditionally certain of her beliefs. Second, *echo*, according to which the unconditional beliefs implied by the condition must be held given the condition. Third, *determination*, which says that the conditional beliefs are the unconditional beliefs that are conditionally certain. Echo implies that conditioning on an event is the same as conditioning on the event being certain, which formalizes the standard informal interpretation of conditional probability. The game-theoretic application of our model, discussed within an example, sheds light on a number of issues in the analysis of extensive form games. Type spaces are closely related to the sphere models of counterfactual conditionals and to models of hypothetical knowledge.

Keywords: Conditional probability; Type spaces; Hypothetical knowledge; Counterfactuals

JEL Classification: C70; C72; D80; D82

*We thank the audiences at various seminar and conference presentations for many helpful comments. We are also grateful to an associate editor and two anonymous referees for many insightful comments. Halpern was supported in part by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, by MURI grant FA9550-12-1-0040, and by ARO grant W911NF-09-1-0281.

[†]IGIER and Department of Economics, Università Bocconi, Italy. alfredo.ditillio@unibocconi.it

[‡]Computer Science Department, Cornell University, U.S.A. halpern@cs.cornell.edu

[§]Faculty of Management, Tel Aviv University, Israel. dovsamet@gmail.com

1 Introduction

The standard models of asymmetric information used in game theory and economics are the type spaces of [Harsanyi \(1967-68\)](#) and the more general partition models of [Aumann \(1976\)](#) and belief spaces of [Mertens and Zamir \(1985\)](#). In these models, the agents' interactive beliefs are described by specifying, at each state in a state space, for each agent, a probability measure on the events in the space, called the agent's *belief type* at the state. In such models it is impossible to formalize the counterfactual probabilistic thinking that is essential for rational choice in extensive form games—for example, a player's assessment of the relative likelihood of continuations of play that follow actions which she is certain not to choose. Probabilistic beliefs are inadequate for describing this kind of thinking, because conditional probabilities are not defined when the conditioning event has probability zero.

The straightforward way of modeling counterfactual beliefs is by taking conditional probability as a primitive notion ([Rényi, 1955](#)) rather than deriving it from probability.¹ Here, one specifies a subfamily of events, called *conditions*, and a family of probability measures, one for each condition, satisfying normality—each is supported by the condition it is associated to—and the chain rule—they are related to each other by Bayes rule whenever possible. In this paper we model beliefs expressed by conditional probabilities analogously to the standard modeling of beliefs by probabilities. Thus, focusing on the case of a single agent, we consider models where each state is associated with a *conditional belief type*, a specification of a probability measure for each nonempty condition in a fixed condition field, satisfying normality and the chain rule. (The straightforward generalization to several agents is discussed in the final section of the paper.)

The received literature offers two models of counterfactual thinking closely related to ours, the type spaces of [Battigalli and Siniscalchi \(1999\)](#) and the model of hypothetical knowledge of [Samet \(1996\)](#).² Roughly speaking, the model discussed in this paper shares the basic logic with the latter and the use of conditional probabilities with the former. Our model can be also viewed as an extension of the models for counterfactual conditionals proposed by [Stalnaker \(1968\)](#) and [Lewis \(1973\)](#). These relationships deserve in-depth analysis and detailed discussion, which we therefore provide in the main body of the paper. In the remainder of this introduction, we summarize our framework and results and provide some motivation and intuition for them.

To describe our analysis, we first recall some previous work on the standard model of a belief type space. The latter model allows us to express interactive beliefs by formalizing any statement of the form “the probability of E is at least p ,” where E is an event and p a number, itself as an event. This event, denoted $B^p(E)$, is the set of states where the belief type assigns probability at least p to E . [Samet \(2000\)](#) studied the correspondence between this model-

¹The idea of taking conditional probability as primitive dates back to [Keynes \(1921\)](#), [Popper \(1934, 1968\)](#) and [de Finetti \(1936\)](#). [Rényi \(1955\)](#) was the first to provide a rigorous measure-theoretic treatment.

²For other, less closely related models, see [Feinberg \(2005\)](#) and the references therein.

theoretic approach to the description of beliefs and a more basic, axiomatic approach. In the latter, for each p , “the probability of ... is at least p ” is formalized by an operator $B^p(\cdot)$ on events that is taken as *primitive* rather than derived from belief types as above. In particular, Samet (2000) provided an axiomatic characterization of the model of a belief type space, by identifying axioms on the operators that are necessary and sufficient for the existence of belief types from which the operators are derived.³ In addition, he considered further axioms that characterize subfamilies of models. For example, he showed that the axiom of *introspection*, which requires certainty of one’s own beliefs, or $B^p(E) \subseteq B^1(B^p(E))$ for every E and p , characterizes Aumann’s partition model, that is, it holds (only) in those belief type spaces where, at each state, the belief type assigns probability one to the set of states with that same type.

In this paper we take a similar approach to the study of a state space with conditional belief types. In such a model we can formalize the statement “the probability of E given C is at least p ,” where E is an event and C a condition, as the event consisting of all the states at which the probability measure associated with C , assigns probability at least p to E . Similarly to the case of probabilistic belief statements, we can alternatively formalize “the probability of E given C is at least p ” using a *binary* operator $B^p(\cdot | \cdot)$ mapping event-condition pairs into events, as the event $B^p(E | C)$. Here, too, the questions arise as to under what axioms the two formalizations are equivalent, and which subfamilies of models correspond to certain desirable additional axioms. Our main goal is to provide answers to these questions.

First, we characterize axiomatically the basic model of a conditional belief type space. We list seven axioms on the operators $B^p(\cdot | \cdot)$ such that a state space Ω with operators satisfying these axioms is necessarily a space with conditional belief types that induce the operators. The axioms are of two kinds. The first five axioms, similar to the ones used in Samet (2000), imply that for each condition C , the operator $B^p(\cdot | C)$ is derived from a mapping assigning a probability measure $\iota^\omega(\cdot | C)$ to each state ω . In particular, this defines a belief type space when we consider the *unconditional type* at each state ω , namely the

³Gaifman (1986) also defined belief spaces (which he named *high order probability spaces*) and characterized them in terms of axioms imposed on an operator that maps each event E and closed interval I into another event, described as “the probability of E lies in I .” The phrase “for agent i the probability of ... is at least p ” can be formalized as an operator in a formal language, rather than a set-theoretic operator. This gives rise to a modal logic of probabilistic beliefs for which type spaces serve as semantical models. The most notable examples are Fagin, Halpern, and Megiddo (1990), Fagin and Halpern (1994) and Heifetz and Mongin (2001). The axioms in such languages are analogous to the axioms on set-theoretic operators. However, the modal logic approach has to overcome problems that arise because the field of real numbers is Archimedean. These problems can be circumvented by using a richer language (Fagin and Halpern, 1994) that allows the description of expectations, by introducing a strong inference rule (Heifetz and Mongin, 2001), or by adopting infinitary propositional language and inference rules (Meier, 2012). The set-theoretic axiomatic approach, which is free of the finitary nature of a formal language, avoids these problems while preserving the appeal of the axioms. See Halpern (1999b) for a comparison between the syntactic and set-theoretic axiomatizations for the logics of knowledge, belief, and counterfactuals.

probability measure $t^\omega(\cdot | \Omega)$, and the associated *unconditional belief* operators $B^p(\cdot | \Omega)$, which we abbreviate as $t^\omega(\cdot)$ and $B^p(\cdot)$, respectively. The second two axioms guarantee that at each state ω the probability measures associated to the different conditions form, in fact, a conditional belief type, so that $t^\omega(\cdot | \cdot)$ satisfies normality and the chain rule.

Next, we examine three additional axioms that imply some structure on the type space. First we consider *introspection*, a common assumption in the modeling of knowledge and belief, which roughly says that the agent has full access to her own mental state. In the context of interactive probabilistic beliefs, introspection means that the agent is certain of (i.e. assigns probability one to) her probabilistic beliefs. For interactive conditional beliefs, introspection says that she is unconditionally certain of her conditional beliefs. In terms of the operators, this idea takes the form of the following axiom: $B^p(E | C) \subseteq B^1(B^p(E | C))$. That is, if the probability of E given C is at least p , then this fact is unconditionally certain. We show, similarly to Samet (2000), that introspection can be equivalently expressed in terms of types as certainty, at each state, of the type at the state. Thus, the axiom of introspection holds in a type space if and only if for each element π in the partition of the state space into events of the same type, π is unconditionally certain at π .

The second axiom, which we call *echo*, relates the conditional belief given a condition to the unconditional belief implied by the condition. Formally, the axiom requires that if $C \subseteq B^p(E)$ then $B^p(E | C) = \Omega$, that is, if a condition C implies the unconditional belief $B^p(E)$ then the conditional belief $B^p(E | C)$ is sure.⁴ We show that echo implies introspection and, moreover, that it implies the equality $C = B^1(C)$ for each condition C . Finally, we prove that a type space satisfies echo if and only if, at each state ω , the conditional probability $t^\omega(E | C)$ is the expectation of the unconditional probabilities $t^{\omega'}(E)$ at the states ω' in C , where the expectation is taken with respect to $t^\omega(\cdot | C)$ itself. A corollary of this equivalence is that $t^\omega(\cdot | C)$ is a convex combination of unconditional types, hence a prior on the state space (Samet, 1998a).

Finally, we consider the axiom of *determination*, which requires conditional beliefs to be conditionally certain to be the unconditional ones. Formally, $B^p(E | C) \subseteq B^1(B^p(E) | C)$. In terms of types, determination posits the tightest link between conditional and unconditional probabilities. Indeed, it implies that at each state ω , the conditional probability given a condition C is the unconditional probability of a single, determined unconditional type in C . In particular, under determination, $t^\omega(\cdot | C)$ assigns positive probability to a single element π of the partition of the state space into events of the same type, and $t^\omega(\cdot | C) = t^{\omega'}(\cdot)$ for any state ω' in π . Thus, by the characterization stated above for echo, determination implies echo.

While the introspection axiom and its characterization in terms of types are analogous to those appearing in Samet (2000), echo and determination deserve some more discussion. To

⁴Various versions of axioms that relate conditional probabilities to unconditional probability have been studied. Some authors call such axioms *Miller's principle* after Miller (1966), who claimed that a certain version of this axiom is paradoxical. See Samet (1999) and the discussion and references therein.

motivate the two axioms and shed light on their characterizations, consider the following two representative quotes, which describe the standard *informal* interpretation of conditioning in probability theory. The first is from Wikipedia:

In probability theory, a conditional probability is the probability that an event will occur, when another event is known to occur or to have occurred.

The second is from Billingsley (1995, p. 427):

It is helpful to consider conditional probability in terms of an observer in possession of partial information. As always, observer, information, know, and so on are informal, nonmathematical terms.

The first quote suggests that we interpret the conditional probability of an event as the unconditional probability assigned to the event under the assumption of knowledge of the condition. However, as the second quote emphasizes, knowledge is not a formal notion in the standard model of a probability space—in that model, knowing or even just being certain of an event is not itself an event, hence the interpretation must remain informal. In a conditional belief type space, knowledge and certainty do have formal expression, and thus the informal interpretation of conditioning can be turned into a mathematical property. The axioms of echo and determination are—to different extents—what delivers this property.

Faithfully to the first quote, under echo, at each state ω the conditional probability of an event *given* a condition C echoes the unconditional probabilities assigned to the event *at* the states in C . By the equality $C = B^1(C)$, which is also implied by echo, these unconditional probabilities are the ones assigned to the event, under the assumption that (i.e. at the states in which) the condition C is known.⁵ There may be multiple such probabilities, as C may be compatible with multiple unconditional beliefs, to wit, the unconditional type $t^{\omega'}(\cdot)$ may vary as ω' varies in C . Thus, for each event E , an expectation of the unconditional probability $t^{\omega'}(E)$ is taken, with $t^{\omega}(\cdot | C)$ serving as a prior. Determination requires, in addition, that this prior be concentrated on a single unconditional probability, thus pushing the formal rendering of “the probability of E given C ” as “*the* probability of E when C is known to occur” to its fullest extent.

The motivating application of our model is to the analysis of games in extensive form. In this paper we formally deal only with an example, but we do lay out and discuss the basic building blocks of a more general and complete analysis, which is the object of ongoing work. We illustrate three main points. First, under an additional property which we call *planning*, a player’s strategy at a state can be *defined* from actual behavior and beliefs, rather than

⁵We do not formally introduce knowledge into our analysis, but this is immaterial for our discussion. The partition into events of the same type can be used to define a knowledge operator $K(\cdot)$ on the state space, and this would be the unique knowledge operator satisfying both $B^1(E) \subseteq K(B^1(E))$ and $K(E) \subseteq B^1(E)$ for every event E . Then, echo would imply that $C = B^1(C) = K(C)$ for every condition C . Thus, under echo, conditioning on C can be also formalized as having knowledge of C . See Halpern, Samet, and Segev (2009) for a treatment on the definition of knowledge in terms of belief.

being assumed. Second, again under planning, our model allows violations of (and hence a formal definition for) the traditional requirement that, in comparing different strategies of hers, a player uses the same belief about the other players' strategies. Third, in our model we can express informational assumptions usually thought of as characteristics of a *game*, such as imperfect information, as properties of *types*, that is, in terms of the informational model, thus relating the two notions of information.

The rest of the paper goes as follows. After giving some preliminaries in Section 2, we axiomatize the basic model of a conditional belief type space in Section 3. In Section 4 we analyze introspection, echo and determination. These are further discussed in Section 5, where we present the game-theoretic example. In Section 6 we provide a comparison with three most closely related models. Finally, in Section 7 we make a few final remarks.

2 Preliminaries

Throughout the paper we fix a triple $(\Omega, \mathcal{F}, \mathcal{C})$, where Ω is a finite set of *states*, $\mathcal{F}$ is a field of subsets of Ω called *events*, and $\mathcal{C}$ is a subfield of $\mathcal{F}$ called the *condition field*. A *condition* is a nonempty event in $\mathcal{C}$, that is, an event in $\mathcal{C}^+ = \mathcal{C} \setminus \{\emptyset\}$.

2.1 Probability and conditional probability

A *probability measure* on $(\Omega, \mathcal{F})$ is a function $P: \mathcal{F} \rightarrow [0, 1]$ satisfying *normality*, that is, $P(\Omega) = 1$, and *additivity*, that is, for all $E, F \in \mathcal{F}$, if $E \cap F = \emptyset$ then $P(E \cup F) = P(E) + P(F)$. The set of all probability measures on $(\Omega, \mathcal{F})$ is denoted $\Delta(\Omega, \mathcal{F})$. Given two events E, C with $P(C) > 0$, we let $P(E | C) = P(E \cap C) / P(C)$ and call it the *probability of E given C*. Clearly, the function $P(\cdot | C)$ so defined is a probability measure on $(\Omega, \mathcal{F})$.

In order to define $P(\cdot | C)$ without the requirement that $P(C) > 0$, we take conditional probability as primitive, rather than deriving it from probability. A *conditional probability measure* on $(\Omega, \mathcal{F}, \mathcal{C})$ is a function $P: \mathcal{F} \times \mathcal{C}^+ \rightarrow [0, 1]$, where we write $P(E | C)$ for $P(E, C)$, satisfying the following properties, for all $E, F \in \mathcal{F}$ and $C, D \in \mathcal{C}^+$:

- (N) $P(C | C) = 1$;
- (A) $P(E \cup F | C) = P(E | C) + P(F | C)$ if $E \cap F = \emptyset$;
- (C) $P(E | C) = P(E | D)P(D | C)$ if $E \subseteq D \subseteq C$.

The set of all conditional probability measures on $(\Omega, \mathcal{F}, \mathcal{C})$ is denoted $\Delta(\Omega, \mathcal{F}, \mathcal{C})$. The (conditional) normality and additivity properties, (N) and (A), ensure that for each condition C the function $P(\cdot | C)$ is a probability measure in $\Delta(\Omega, \mathcal{F})$, one putting probability 1 on C . The probability measure $P(\cdot | \Omega)$ is called the *unconditional part* of P , and we often omit the condition, writing $P(\cdot)$. Property (C), the *chain rule*, imposes some relations between

the probability measures. In particular, it follows from (N), (A) and (C) that if $P(C) > 0$ for a condition C , then $P(E | C) = P(E \cap C)/P(C)$ for each event E .⁶

2.2 Lexicographic systems

A conditional probability measure can be equivalently described by a *lexicographic system* on $(\Omega, \mathcal{F}, \mathcal{C})$; we make use of this equivalence in the rest of the paper. A lexicographic system consists of two sequences, $(S_1, \dots, S_k)$ and $(P_1, \dots, P_k)$. The first sequence is an ordered partition of Ω , the elements of which are conditions, namely, events in $\mathcal{C}^+$. The second sequence consists of probability measures on Ω such that for each i , $P_i(S_i) = 1$, and S_i does not contain any condition C such that $P_i(C) = 1$. The sequence $(S_1, \dots, S_k)$ is a hierarchy of conditions such that each S_i is infinitely more likely than S_j for $j > i$. Thus, S_1 is infinitely more likely than its complement $\Omega \setminus S_1$, S_2 is a subset of $\Omega \setminus S_1$ which is infinitely more likely than its complement in $\Omega \setminus S_1$, namely $\Omega \setminus (S_1 \cup S_2)$, and so on. For any event E in $\mathcal{F}$ and index i , the probability $P_i(E)$ is the likelihood of E conditional on S_i , that is, at the i th level of the hierarchy.

There is a one-to-one correspondence between conditional probability measures and lexicographic systems, which we describe next. Starting with a conditional probability P we construct a lexicographic system, $(S_1, \dots, S_k)$ and $(P_1, \dots, P_k)$, as follows. The event S_1 is the smallest event in $\mathcal{C}$ that is unconditionally certain,⁷ and $P_1(\cdot)$ is just $P(\cdot | \Omega)$. Since S_1 is a condition, so is its complement $\Omega \setminus S_1$, and moreover this event is unconditionally null. However, $P(\cdot | \Omega \setminus S_1)$ is a probability over this unconditionally null set, and S_2 is the smallest event in $\mathcal{C}$ that is conditionally certain given $\Omega \setminus S_1$. The probability measure $P_2(\cdot)$ is $P(\cdot | \Omega \setminus S_1)$, and so on. Formally, starting with $S_0 = \emptyset$, for each $i > 0$ we let $P_i(\cdot) = P(\cdot | \Omega \setminus (S_0 \cup \dots \cup S_{i-1}))$, and define S_i as the minimal set in $\mathcal{C}$ for which $P_i(S_i) = 1$. Clearly, the sets S_i thus defined are disjoint and nonempty, and therefore, for some $k \geq 1$, the sequence $(S_1, \dots, S_k)$ is an ordered partition of Ω . The minimality of each S_i guarantees that S_i does not contain any condition C such $P_i(C) = 1$. Thus the two sequences form a lexicographic system.

Conversely, starting from a lexicographic system $(S_1, \dots, S_k)$ and $(P_1, \dots, P_k)$, for each event E and condition C we define the conditional probability of E given C as the likelihood of E within the highest-level condition S_i compatible with C . Formally, for each condition C we let $P(\cdot | C) = P_i(\cdot)$, where i is the smallest index such $C \cap S_i \neq \emptyset$. By the minimality of S_i and our assumption that $\mathcal{C}$ is a field, which implies that $S_i \cap C$ is a condition, and also that $S_i \cap D$ is a condition for each condition D with $D \subseteq C$, it easily follows that P so defined is indeed a conditional probability measure. Moreover, the lexicographic system

⁶Myerson (1986b, pp. 336–337, and 1986a) defines a conditional probability measure for the case $\mathcal{C} = \mathcal{F}$. Variants of conditional probabilities are also studied by Hammond (1994) and Halpern (2010).

⁷When $\mathcal{C} = \mathcal{F}$, then S_1 is simply the support of $P(\cdot | \Omega)$.

associated above with P is $(S_1, \dots, S_k)$ and $(P_1, \dots, P_k)$.⁸

Given a conditional probability measure P , we refer to the sequence $(S_1, \dots, S_k)$ in the corresponding lexicographic system as the *hierarchy* induced by P . Given a condition C , we call the condition $C^+ = C \cap S_i$, where i is the smallest index for which $C \cap S_i \neq \emptyset$, the P -positive part of C . It follows immediately from the description of the correspondence that the only part of C that matters for conditioning is its P -positive part C^+ , that is:

Claim 1. *For each condition C , $P(\cdot | C) = P(\cdot | C^+)$.*

3 Conditional belief types

3.1 Type functions and belief operators

In order to express statements about conditional beliefs as events, we consider a state space where each state is associated with conditional beliefs on the state space, much as in a standard belief space unconditional beliefs are associated with states. Here, a *type function* is a function $t : \Omega \rightarrow \Delta(\Omega, \mathcal{F}, \mathcal{C})$ which assigns to each state a conditional probability measure on $(\Omega, \mathcal{F}, \mathcal{C})$. For each event E and condition C , the function $t(\cdot)(E | C)$ is required to be measurable with respect to $\mathcal{C}$. That is, for each $p \in [0, 1]$,

$$\{\omega \in \Omega : t(\omega)(E | C) \geq p\} \in \mathcal{C}.$$

This measurability condition, which is stronger than measurability with respect to $\mathcal{F}$, enables conditioning on the events concerning the agent's conditional beliefs themselves. In all our results up to and including Corollary 2, with the exception of Corollary 1, the condition can be entirely dispensed with. For the remaining results, it can be replaced with the weaker requirement that events concerning conditional certainty be in the condition field. That is, for each event E , the set $\{\omega \in \Omega : t(\omega)(E) = 1\}$ is in $\mathcal{C}$. This is because under the axiom of introspection, to be introduced later, the two conditions turn out to be equivalent.

In what follows, for each state ω we write t^ω for $t(\omega)$, and call it the *type at ω* . We also write $t^\omega(\cdot)$ instead of $t^\omega(\cdot | \Omega)$ and we call it the *unconditional type at ω* . Obviously, the space $(\Omega, \mathcal{F})$ and the function $\omega \mapsto t^\omega(\cdot)$ define an unconditional probability type space.

A *family of conditional belief operators* (a *family of operators*, for short) is a collection of operators $(B^p)_{p \in [0,1]}$ where $B^p : \mathcal{F} \times \mathcal{C}^+ \rightarrow \mathcal{C}$ for each $p \in [0, 1]$. For an event E and

⁸Rényi (1956) describes an equivalence relation on conditions which in the finite case results in the ordered partition described here. He further defines *dimensionally ordered measures* which in the finite case are the probability measures P_i . Lexicographic systems of conditional probability, for $\mathcal{C} = \mathcal{F}$, are studied in Blume, Brandenburger, and Dekel (1991, pp. 71–72). A proof of the equivalence between the axiomatic description of conditional probabilities and their description in terms of lexicographic systems, for the case $\mathcal{C} = \mathcal{F}$, appears in Monderer, Samet, and Shapley (1992). Here, conditional probabilities are presented for the more general case where $\mathcal{C}$ is any subfield of $\mathcal{F}$.

condition C , we write $B^p(E | C)$ rather than $B^p(E, C)$. It is the event that the *belief in E given C is at least p* . If $C = \Omega$, we omit the condition and write just $B^p(E)$. This is the event that the *unconditional belief in E is at least p* . The requirement that the images of the operators B^p are in $\mathcal{C}$, rather than $\mathcal{F}$, is imposed in order to enable conditioning on events concerning beliefs. This is analogous to the measurability condition on the type function t , and the analogous remarks apply here—in particular, for our purposes it suffices to assume that the image of the unconditional belief operator $B^1(\cdot)$ is in $\mathcal{C}$.

A type function t corresponds in a natural way to a family of operators: the event that the belief in E given C is at least p consists of all the states where the type assigns a probability of at least p to E given C . Formally, for all $E \in \mathcal{F}$, $C \in \mathcal{C}^+$, and $p \in [0, 1]$, we let

$$(1) \quad B^p(E | C) = \{\omega \in \Omega : t^\omega(E | C) \geq p\}.$$

Now we introduce axioms that characterize the families of operators that correspond to type functions. For all $E, F \in \mathcal{F}$, $C, D \in \mathcal{C}^+$, and $p, q, p_n \in [0, 1]$:

- (P1) $B^0(E | C) = \Omega$;
- (P2) $B^p(E \cap F | C) \cap B^q(E \cap \neg F | C) \subseteq B^{p+q}(E | C)$ for $p + q \leq 1$;
- (P3) $\neg B^p(E \cap F | C) \cap \neg B^q(E \cap \neg F | C) \subseteq \neg B^{p+q}(E | C)$ for $p + q \leq 1$;
- (P4) $B^p(E | C) \cap B^q(\neg E | C) = \emptyset$ for $p + q > 1$;
- (P5) $\bigcap_n B^{p_n}(E | C) \subseteq B^p(E | C)$ for $p_n \uparrow p$;⁹
- (PN) $B^1(C | C) = \Omega$;
- (PC) $B^p(E | D) \cap B^q(D | C) \subseteq B^{pq}(E | C)$ for $E \subseteq D \subseteq C$.

Axioms (P1)–(P5) and (PN) correspond to the requirement that for each condition C the function $t^\omega(\cdot | C)$ is a probability measure for each ω . Analogous axioms were introduced by Samet (2000, p. 174) for unconditional belief operators, and we refer the reader to that article for a discussion. Axiom (PN) corresponds to the axiom of conditional normality, and axiom (PC) is the counterpart of the chain rule of conditional probability measures. These axioms characterize the families of operators that correspond to type functions.

Theorem 1. *A family of operators corresponds to a type function if and only if it satisfies axioms (P1)–(P5), (PN), and (PC). In this case, the type function is unique.*

The proof of this theorem is in the Appendix.

In the remainder of the paper we fix a conditional type function t and the corresponding family of conditional belief operators (B^p) defined by (1). By Theorem 1, the family (B^p) must satisfy (P1)–(P5), (PN), and (PC). For each state ω , we denote by $(S_1^\omega, \dots, S_k^\omega)$ the hierarchy associated with the conditional probability measure t^ω .

⁹Here $p_n \uparrow p$ means that the sequence $p_1, p_2, \dots$ converges to p from below.

3.2 The belief field

The basic events concerning the agent's conditional beliefs are the ones in the range of the conditional belief operators, namely, the events in the family

$$\mathcal{B} = \{B^p(E \mid C) : p \in [0, 1], E \in \mathcal{F}, C \in \mathcal{C}^+\}.$$

We call the field generated by this set the *belief field*, and denote it by $\mathcal{E}$. Thus, the events in the field $\mathcal{E}$ are those events that describe statements about conditional beliefs. Note that since we required that the images of the belief operators be in $\mathcal{C}$, and since we assumed that $\mathcal{C}$ itself is a field, it follows that $\mathcal{E} \subseteq \mathcal{C}$. Thus, every nonempty event concerning the agent's beliefs is in particular a condition.

The belief field can be equivalently described in terms of the type function that corresponds to the family of conditional belief operators. For the latter description, let Π denote the partition of Ω into states with the same type, so that for each state ω , the element of the partition containing ω is

$$\Pi(\omega) = \{\omega' \in \Omega : t^{\omega'} = t^\omega\}.$$

The next result shows the equivalence between the two descriptions.

Proposition 1. *The field generated by Π is the belief field $\mathcal{E}$.*

The proof of the proposition is in the Appendix.

Proposition 1 implies that, for each state ω , the partition element $\Pi(\omega)$ is an element of the belief field $\mathcal{E}$ and hence an element of $\mathcal{F}$, that is, an event—the event that the agent's type is t^ω . Moreover, since $\mathcal{E} \subseteq \mathcal{C}$, it follows that all the elements of the partition Π , and hence all unions of elements of Π , are conditions.

4 Introspection, echo, and determination

Although the family of operators is able to express conditional beliefs about conditional beliefs, the axioms considered so far, (P1)–(P5), (PC), and (PN), make no special provision regarding iterations of the operators, that is, consideration of events $B^p(E \mid C)$ where E or C are themselves events that describe beliefs. This is reflected in the fact that, except for measurability of the type function, no restriction is imposed on how types in different states are related to each other. In this section we introduce three such requirements, expressed in terms of axioms on the family of operators. We study how these axioms are related to each other and investigate their impact on the relationship between types at different states.

4.1 Introspection

Beliefs, conditional or unconditional, are in the agent's mind. The agent satisfies *introspection* if she is unconditionally certain of her beliefs. We formalize this in terms of the belief

operators by the following axiom. For all $E \in \mathcal{F}$, $C \in \mathcal{C}^+$, and $p \in [0, 1]$,

$$(Int) \quad B^p(E | C) \subseteq B^1(B^p(E | C)).$$

Introspection can be equivalently expressed in terms of properties of the type function:

Proposition 2. *Axiom (Int) holds if and only if for each ω , $t^\omega(\Pi(\omega)) = 1$.*

Proof. Since $\mathcal{B}$ is finite, Lemma 2 (in the Appendix) implies that $t^\omega(\Pi(\omega)) = 1$ for each $\omega \in \Omega$ if and only if, for each $\omega \in \Omega$ and $B \in \mathcal{B}$, if $\omega \in B$ then $t^\omega(B) = 1$. This is true if and only if $\omega \in B^1(B)$ for each $B \in \mathcal{B}$ and $\omega \in B$, that is, if and only if (Int) holds. ■

Since $\Pi \subseteq \mathcal{E} \subseteq \mathcal{C}$, and since S_1^ω is the minimal event in $\mathcal{C}$ which is certain for the probability measure $t^\omega(\cdot)$, we obtain from Proposition 2 the following:

Corollary 1. *Axiom (Int) holds if and only if for each ω , $S_1^\omega \subseteq \Pi(\omega)$.*

Introspection can also be expressed in terms of the belief field:

Proposition 3. *Axiom (Int) holds if and only if for each E in the belief field $\mathcal{E}$, $E = B^1(E)$.*

Proof. Axiom (Int) follows from the condition in the proposition by substituting $B^p(E | C)$ for E . Suppose that (Int) holds and let $E \in \mathcal{E}$. Consider a state $\omega \in E$. Since Π generates $\mathcal{E}$, by Proposition 1, E is a union of elements of Π , and thus $\Pi(\omega) \subseteq E$. Therefore, $t^\omega(E) \geq t^\omega(\Pi(\omega)) = 1$, and thus $\omega \in B^1(E)$. If $\omega \in \neg E$, then by the same argument, $t^\omega(\neg E) = 1$. Thus, $t^\omega(E) = 0$, and $\omega \in \neg B^1(E)$. ■

When belief and knowledge are studied, axioms like (Int) are said to capture *positive* introspection. In contrast, *negative* introspection refers to knowing that one does *not* know and believing that one does not believe. For knowledge and belief, negative introspection does not follow from positive introspection.¹⁰ But when probabilistic beliefs are involved, negative introspection is implied by positive introspection. Indeed, since the events of the form $\neg B^p(E | C)$ are in $\mathcal{E}$, negative introspection follows immediately from Proposition 3:

Corollary 2. *If axiom (Int) holds, then for all $E \in \mathcal{F}$, $C \in \mathcal{C}^+$, and $p \in [0, 1]$,*

$$\neg B^p(E | C) = B^1(\neg B^p(E | C)).^{11}$$

¹⁰In modal logic, positive introspection is known as axiom (4) and negative introspection as axiom (5).

¹¹Just like positive introspection, negative introspection is usually expressed with inclusion, that is, in our framework, as $\neg B^p(E | C) \subseteq B^1(\neg B^p(E | C))$. Axiom (P4), which corresponds to axiom (D) in modal logic, makes both positive and negative introspection hold with equality.

4.2 Echo

The next axiom says that if the agent *unconditionally* believes an event E with probability at least p when a condition C holds, then the agent must believe E with probability at least p given the condition C . Formally, for all $E \in \mathcal{F}$, $C \in \mathcal{C}^+$ and $p \in [0, 1]$,

$$(\text{Echo}) \quad \text{if } C \subseteq B^p(E) \text{ then } B^p(E | C) = \Omega.$$

Towards our characterization of echo in terms of types, we explore first the case where C is an element of the partition Π , that is, when the condition is a single type. In this case (Echo) implies that the conditional probability given the type is the unconditional probability of that type. By the definition of Π , for each $\pi \in \Pi$ we can write t^π to denote the type that is constant across all the states in π . Then, we have:

Proposition 4. *If axiom (Echo) holds, then for all $\omega \in \Omega$ and $\pi \in \Pi$,*

$$(2) \quad t^\omega(\cdot | \pi) = t^\pi(\cdot).$$

Proof. Suppose (Echo) holds. If $t^\pi(E) \geq p$ then $\pi \subseteq B^p(E)$. Thus, by (Echo), $t^\omega(E | \pi) \geq p$. As this is true for all p and E , the probability measures $t^\omega(\cdot | \pi)$ and $t^\pi(\cdot)$ coincide. ■

Axiom (Echo) has two important implications.

Proposition 5. *Axiom (Echo) implies both of the following:*

- (a) Axiom (Int).
- (b) $\mathcal{C} = \mathcal{E}$, that is, the condition field and the belief field coincide.

Proof. For each $\omega \in \Omega$, by (2) and (N), $t^\omega(\Pi(\omega)) = t^\omega(\Pi(\omega) | \Pi(\omega)) = 1$. Thus, (Int) follows from Proposition 2.

Now suppose that (Echo) holds, but $\mathcal{C} \neq \mathcal{E}$. Since $\mathcal{E} \subseteq \mathcal{C}$ and Π generates $\mathcal{E}$, there must exist $\pi \in \Pi$ and C' and C'' in $\mathcal{C}^+$ such that $C' \cup C'' = \pi$. Now, if $t^\pi(E) \geq p$ then $C' \subseteq B^p(E)$ and $C'' \subseteq B^p(E)$ and hence, by (Echo), $t^\pi(E | C') \geq p$ and $t^\pi(E | C'') \geq p$. Since this is true for all E and p , it follows that $t^\pi(\cdot | C') = t^\pi(\cdot | C'') = t^\pi(\cdot)$. Thus, $t^\pi(C') = t^\pi(C' | C') = 1$ and $t^\pi(C'') = t^\pi(C'' | C'') = 1$, which is a contradiction, since $C' \cap C'' = \emptyset$. ■

When (Echo) holds, then (Int) holds by Proposition 5. This implies, by Proposition 3, that for each $E \in \mathcal{E}$, $E = B^1(E)$. Finally, each condition C is in $\mathcal{E}$, again by Proposition 5. Thus we conclude:

Corollary 3. *If axiom (Echo) holds, then for each condition C , $C = B^1(C)$.*

Thus, with (Echo), conditioning on C means conditioning on C being unconditionally certain. This is a formalization of the common informal idea that conditional probability is probability under knowledge of the condition (see footnote 5).

The following equivalence theorem extends Proposition 4 for conditioning events that are not a single type, and provides a necessary and sufficient condition for (Echo) in terms of the type function. The condition is that the probability given C at a state is a convex combination of the unconditional types at C with weights that are given by the conditional probability of the types. We discuss this condition in more detail below.

Theorem 2. *Axiom (Echo) holds if and only if for each state ω and condition C ,*

$$(3) \quad t^\omega(\cdot | C) = \sum_{\pi \subseteq C} t^\omega(\pi | C) t^\pi(\cdot).$$

Proof. Suppose (Echo) holds. By part (b) of Proposition 5, each $C \in \mathcal{C}^+$ is a union of elements of Π . By normality and additivity, for each ω , E , and C , $t^\omega(E | C) = \sum_{\pi \subseteq C} t^\omega(E \cap \pi | C)$. Applying the chain rule to each summand, then normality, and finally (2), we obtain:

$$t^\omega(E \cap \pi | C) = t^\omega(E \cap \pi | \pi) t^\omega(\pi | C) = t^\omega(E | \pi) t^\omega(\pi | C) = t^\pi(E) t^\omega(\pi | C).$$

Since this holds for each E , (3) follows. Conversely, suppose that (3) holds. Then $\mathcal{C} = \mathcal{E}$. Indeed, if this were not the case, then, as in the proof of part (b) in Proposition 5, there is a condition C which is a nontrivial subset of some $\pi \in \Pi$. For such C the sum righthand side of (3) has no summands and the equation cannot hold. Thus, each condition is the union of elements of Π . Assume that for a condition C , $C \subseteq B^p(E)$. For $\pi \subseteq C$, $\pi \subseteq B^p(E)$ and hence $t^\pi(E) \geq p$. Thus, by (3), for any ω , $t^\omega(E | C) = \sum_{\pi \subseteq C} t^\pi(E) t^\omega(\pi | C) \geq p \sum_{\pi \subseteq C} t^\omega(\pi | C) = p$. Thus for each ω , $\omega \in B^p(E | C)$, that is, (Echo) holds. ■

The summation in (3) can be taken for elements π in a subset of C , as we state next.

Corollary 4. *Axiom (Echo) holds if and only if for each state ω and condition C ,*

$$(4) \quad t^\omega(\cdot | C) = \sum_{\pi \subseteq C^+} t^\omega(\pi | C) t^\pi(\cdot),$$

where C^+ is the t^ω -positive part of C .

Indeed, plug C^+ for C in (3), and then replace $t^\omega(\cdot | C^+)$ with $t^\omega(\cdot | C)$, using Claim 1.

4.2.1 Conditional probabilities as priors

It is well known that in (unconditional) belief spaces that satisfy introspection, a probability measure on the space is a prior for an agent if and only if it is a convex combination of the agent's types (see Samet, 1998a). Moreover, in such spaces, a probability measure is a

prior if and only if it is an invariant probability of the Markov chain on the type space where the types at the states are the transition probability measures (see Samet, 1998b). Now, equation (3) shows that the conditional probability given C is a convex combination of the unconditional types. Moreover, this equation can be equivalently written as,

$$t^\omega(\cdot | C) = \sum_{\pi \subseteq C} t^\omega(\pi | C) t^\pi(\cdot) = \sum_{\pi \subseteq C} \sum_{\omega' \in \pi} t^\omega(\omega' | C) t^{\omega'}(\cdot) = \sum_{\omega' \in C} t^\omega(\omega' | C) t^{\omega'}(\cdot).$$

Thus, the conditional probability $t^\omega(\cdot | C)$ is an invariant probability of the Markov chain with transition probability measures that are the unconditional probability measures $t^{\omega'}(\cdot)$. We conclude:

Corollary 5. *If axiom (Echo) holds, then for each ω and C , the probability measure $t^\omega(\cdot | C)$ is a prior of the unconditional type space, and in particular it is an invariant probability of the Markov chain whose transition probabilities are given by the unconditional types $t^{\omega'}(\cdot)$.*

4.3 Determination

Like (Echo), the next axiom, which we call *determination*, relates conditional beliefs to unconditional beliefs: for all $E \in \mathcal{F}$, $C \in \mathcal{C}^+$ and $p \in [0, 1]$,

$$(\text{Det}) \quad B^p(E | C) \subseteq B^1(B^p(E) | C).$$

Unlike (Echo), the unconditional beliefs here are not those held *at* the condition, but rather the unconditional beliefs that are *conditionally certain*. When the agent assigns to E a probability of at least p given C , she is certain that this is her unconditional belief, given C .

Axiom (Det) turns out to be stronger than (Echo). Therefore, when (Det) holds, for each ω and C , $t^\omega(\cdot | C)$ is a convex combination of the unconditional types at C . However, with (Det), this convex combination is trivial, and consists of a single type at C . Thus, beliefs given C are the beliefs of a *determined* type in C . As the following theorem states, this determined type is the most probable one in C with respect to the type t^ω .

Theorem 3. *Axiom (Det) holds if and only if the following two conditions hold:*

- Axiom (Echo) is satisfied;
- for each state ω , the hierarchy $(S_1^\omega, \dots, S_{k^\omega}^\omega)$ consists of elements of the partition Π .

Thus, if axiom (Det) holds, then for each ω and C , $t^\omega(\cdot | C) = t^\pi(\cdot)$, where π is the t^ω -positive part of C .


Proof. Suppose that (Det) holds, and let $C \subseteq B^p(E)$. Assume that contrary to (Echo) there exists $\omega \notin B^p(E | C)$. This implies that for some $q > 1 - p$, $\omega \in B^q(\neg E | C)$. Thus, by (Det), $\omega \in B^1(B^q(\neg E) | C)$. Therefore $C \cap B^q(\neg E) \neq \emptyset$, and as $C \subseteq B^p(E)$,

$B^p(E) \cap B^q(\neg E) \neq \emptyset$ which is impossible. Now consider an element S_i^ω of the hierarchy associated with t^ω . Assume $\pi \subseteq S_i^\omega$. By the definition of S_i^ω , $t^\omega(\pi \mid \Omega \setminus (S_1^\omega \cup \dots \cup S_{i-1}^\omega)) > 0$, and hence, $t^\omega(\pi \mid S_i^\omega) > 0$. Now, if $t^\omega(E \mid S_i^\omega) \geq p$, then $t^\omega(B^p(E) \mid S_i^\omega) = 1$, by (Det). Thus, $\pi \cap B^p(E) \neq \emptyset$, which implies by the definition of Π , $\pi \subseteq B^p(E)$. Thus, $t^\pi(E) \geq p$. Since this is true for each E and p , it follows that $t^\omega(\cdot \mid S_i^\omega) = t^\pi(\cdot)$. If $\pi' \neq \pi$, then $t^\omega \neq t^{\pi'}$, and thus, π' is not a subset of S_i^ω .


Suppose now that the two conditions in the theorem hold. Then, for each ω and C , the t^ω -positive part of C is an element of Π . Therefore, by Corollary 4, for each ω , E , and C , $t^\omega(E \mid C) = t^\pi(E)$, where π is the t^ω -positive part of C . Note also that by Proposition 5 and Proposition 2, $t^\pi(\pi) = 1$. Now, to prove that (Det) holds, assume that $t^\omega(E \mid C) \geq p$. Then, $t^\pi(E) \geq p$. Hence, $\pi \subseteq B^p(E)$. But, $t^\omega(\pi \mid C) = t^\pi(\pi) = 1$, so $t^\omega(B^p(E) \mid C) = 1$. ■

5 Game-theoretic application


In this section we use an example to illustrate how the model developed in the previous sections can be used for game-theoretic analysis. Figure 1 represents a two-player game in extensive form and a model of conditional belief types, representing the beliefs the players entertain *before* the game is played, as in Samet (1996). Similarly to the latter paper, we assign to each state a play path, not a strategy, which will instead be derived from the player's conditional beliefs.¹² The state space has six states, each indexed by a subscript indicating the last action in the assigned path. We arrange the states in a matrix which, when partitioned into rows (resp. columns), represents player 1's (resp. player 2's) partition into states of the same type. For convenience, we represent the conditional type functions of the two players in two different diagrams—the one in the middle for player 1, the one on the right for player 2.



The game.



Player 1's type function.



Player 2's type function.

Figure 1: A model for a game in extensive form.

The numbers appearing inside a player's type (partition element) are the unconditional

¹²See also Battigalli, Di Tillio, and Samet (2013), discussed later on, for a similar approach.

probabilities assigned by the type to the states in it. The number attached to the arrow going from a type to another type is the conditional probability assigned by the first type to the second type, given the condition that the type is not the first type. As we require echo and each player has at most three types, this is enough to describe the type functions. Thus, for example, the first type of player 1 assigns unconditional probability 0.5 to state ω_a , conditional probability 0.3 to state ω_e given the condition that her type is the second type, conditional probability 0.7α to state ω_g given the condition that her type is either her second or her third type, and so on. Note that determination holds if and only if each of α , β and γ is either zero or one, so that each type assigns conditional probability one to a single other type, given the condition that the type is not the first type.

In what follows, we identify each node in the game with the last action leading to it. Thus, for instance, we identify the sequence of actions (b, c) with action c. For each action x in the game, we write $[x]$ to denote the set of states where x is chosen in the path at the state. Thus, for instance, $[a] = \{\omega_a, \omega'_a\}$, $[c] = \{\omega_e, \omega_f\}$, and so on.

5.1 Defining strategies from types

In our example, each player is certain of the actions she takes, but her type is compatible with every action of the other player which she (the first player) does not herself exclude. Consider, for example, state ω_g . In this state player 1 is unconditionally certain of her action (b) at the initial node, and also of her actions (e and g) at the two nodes of hers that are not excluded by her initial choice (nodes c and d, respectively), both of which her type allows. Similarly, the type of player 2 in state ω_h is compatible with both initial actions by player 1, assigns unconditional probability 1 to the event that d is chosen if b is chosen, and is compatible with both g and h, the actions of player 1 that are not excluded by d. Formally, our example satisfies the following. First, for every node v and action x at v , if the event $\neg[v] \cup [x]$ is true, then it is given probability 1 by the player moving at v . Second, for each action x of a player, no event in the belief field of the *other* player implies the event $[x]$.

It can be readily verified that the properties illustrated in the previous paragraph, which we jointly refer to as *planning*, can be equivalently stated as follows. First, the type of a player determines her actions, namely, for any given node of hers, she takes the same action in all states of her type, at which the node is reached. Second, the type of a player does not determine the other player's actions, that is, if a node of the other player is reached in some state of the first player's type, then each action at the node is taken in some state of that type. These two properties imply that each state is associated with a unique reduced strategy—by which we mean a strategy in the reduced normal form of the game—for each player. Moreover, the assignment of a player's reduced strategies to states is measurable with respect to the partition of the space into the player's types.¹³ For instance, in state ω_g

¹³Thus, in terms of the knowledge operator induced by this partition, either the player knows she takes an

the profile of reduced strategies is (beg, d), in state ω_f it is (bfh, c), in state ω_a it is (a, c), and so on.¹⁴

Note that planning is independent of the axioms of echo and determination. Planning may hold in a model where echo (and a fortiori determination) fails, and fail in a model where determination (and a fortiori echo) holds. The reason is that planning imposes restrictions on a type only in terms of events (reached nodes) that do not contradict the type, while echo and determination—as opposed to introspection alone—put restrictions on a type only in terms of events (conditions) that are incompatible with (disjoint from) the type itself. Indeed, note that under introspection, whether planning holds or not can be checked by simply considering the certain events, because (Proposition 3) the phrase “no event in the belief field of the other player” can be replaced by “no event that the other player is certain of”. In particular, planning only requires the holding or failing of certain implications between such events and events concerning the game play (nodes being reached and actions being taken).

While reduced strategies are definable at each state in the example, extending this procedure to define *strategies* at each state is possible if and only if α is either 0 or 1, which is the case if, in particular, determination holds. Indeed, in states ω_a and ω'_a player 1 chooses a, but given that she does not, that is, conditional on [b], she assigns probabilities α and $1 - \alpha$ to her other two types. Thus, if $0 < \alpha < 1$, then given [b] she assigns positive probability to two distinct continuation strategies (eg and fh) for the subgame that follows b. In other words, her plan is not *deterministic* for the part of the tree that she herself excludes, but (by echo) is rather a mixture of deterministic plans. If, instead, $\alpha = 0$ or $\alpha = 1$, then in states ω_a and ω'_a the probability measure associated with condition [b] equals the unconditional part of her third or second type, respectively. This defines a strategy, not just a reduced strategy, for her first type. This strategy is afh if $\alpha = 0$ and aeg if $\alpha = 1$.

In the discussion above we take no stand on whether reduced strategies should be definable at each state, or whether the stronger assumption (guaranteed by determination) concerning strategies should be satisfied. Instead, our main purpose is to illustrate the fact that in our model we can formulate such assumptions explicitly. It should be noted, however, that the assumption of planning is what enables a meaningful analysis of rational play under echo, and in particular under Corollary 3. Indeed, planning lets us formalize a statement like “c and d are equally likely given b,” by allowing us to identify it with “c and d are equally likely given that b is certain.” (This statement is true for the first type of player 1, that is, in states ω_a and ω'_a , if $\alpha = 0.6$.)

action that excludes a node, or she considers the node possible—she does not know the node is not reached—and knows the action she takes there.

¹⁴Not all reduced strategies appear in the model—there is no state where the strategy of player 1 is beh or bfg. This is just an artifact of the limited number of states in the example. If planning holds and there are sufficiently many states, then all reduced strategies can find expression in the model. For example, by introducing two more states, with respective paths e and h, forming a fourth type of player 1, we would include strategy beh. We return to this point later, when we discuss imperfect information.

5.2 Rationality and independence

In the game-theoretic application of our model, rationality is a property of actual behavior. A player is rational if her unconditional expected payoff is no less than her conditional expected payoff, when conditioned on any reduced strategy in the model. Note that planning guarantees that the latter condition is an event concerning the beliefs of the player, namely, the union of all types who choose the reduced strategy under consideration. Thus, in order to compare the chosen strategy with another one, a player considers the types of hers who *actually* choose the second strategy, and assesses her payoff in those situations as her *expected* expected payoff.

It is straightforward to see that only the first type of player 1 behaves rationally in the example above. She chooses **a** and her unconditional expected payoff is 2. By echo, her conditional expected payoff given that she chooses **beg** is the unconditional expected payoff of her second type, namely $0.3 \times 5 = 1.5$. Similarly, her conditional expected payoff given that she chooses **bfb** is the unconditional expected payoff of her third type, which is $0.2 \times 6 = 1.2$. Finally, for both her second and third type, the conditional expected payoff given **a** is 2, which makes her behavior irrational.¹⁵

Inherent to the definition of rationality in traditional models, is an assumption of *independence* between a player's choices and beliefs. Rationality of a player is determined by comparing her actual expected payoff to her expected payoff when her strategy is changed, while her beliefs—and in particular her belief about the other players' choices—are kept fixed. By contrast, in our model rationality is checked by conditioning on unchosen reduced strategies, and it is possible that under such conditions, beliefs about other players' reduced strategies change as well. Thus, independence must be stated explicitly: a player's unconditional distribution over the other players' reduced strategies is the same as her conditional distribution, when conditioned on any of her reduced strategies.

In our example, independence fails at every state for both players. For instance, the first type of player 1, who chooses **a**, assigns equal unconditional probabilities to **c** and **d**, but conditional probability 0.3 to **c**, given **beg**.¹⁶ Similarly, the first type of player 2, who plans to choose **c** after **b**, expects a payoff of 0.8, but a payoff of 0 when conditioning on choosing **d**. Note that in the reduced strategic form of the game, **a** is never a best response for player 1,

¹⁵These conclusions hold regardless of the values of α , β and γ , that is, whether echo or determination holds. In our example, each reduced strategy is chosen by at most one type, hence conditioning on a reduced strategy is the same as conditioning on a type. This precludes any behavioral differences between the two axioms. However, in a richer model where multiple types of a player choose the same reduced strategy, such differentiation *can* occur. For example, with echo (and not determination) a type's actual strategy can give a payoff equal to the expected value of the payoff expected by those types who choose another strategy, while none of these expected payoffs is individually equal to it. In this scenario, determination could not deliver such indifference.

¹⁶When each reduced strategy is chosen by at most one type, as is the case in our example, independence requires each type of a player to have the same unconditional belief about the other players' types.

while c and d give player 2 the same payoff for each choice by player 1. It is the violation of independence that allows the dominated strategy a to be rational for player 1 in states ω_a and ω'_a , and the strategies c and d to be perceived differently by player 2.

Finally, it is easy to verify that if independence holds at every state, then rationality has the following intuitive, self-referential characterization. A player is rational if and only if, given any reduced strategy other than the actual one, either her conditional expected payoff is the same as her unconditional one, or the conditional probability of the player being rational is low enough. That is, a rational player who considers reasons to behave differently, conditions on such different behavior, and finds that she either gets the same expected payoff, or behaves irrationally in “too many” such situations.

5.3 Imperfect information

In the received literature, perfect and imperfect information are properties of a game, which determine the players’ strategy sets. Thus, type spaces are models of information concerning a primitive object, the strategies, which already encapsulates a notion of information. The latter kind of information is not formalized in terms of the informational model, and there is no interplay between the two notions of information. In our model, imperfect information can be described in terms of the information already formalized by the type space. The strategies that we derived, as discussed in the previous section, *define* the information structure of the game. Rather than starting with fixed informational assumptions about the game and then formulating a model for it, one can express those assumptions from certain properties of types. Thus, the imperfect information of the *game* is expressed as imperfect information of a *type*.

To illustrate, consider the imperfect-information version of the game of Figure 1, in which player 1 cannot distinguish nodes c and d, and views e and g as being the same action (and similarly for f and h). Clearly, we can view the type functions in the figure as depicting precisely this scenario. Indeed, this is an equally legitimate reading of the example, and there seems to be no reason to expect our predictions to depend on which of the two we may have in mind. The second type of player 1, who chooses beg, has imperfect information, as she cannot imagine keeping b and changing e to f, without also changing g to h. A similar argument applies to her third type. These conclusions are valid, whether we see the state space as a model for the perfect- or imperfect-information version of the game. Of course, in this example player 1’s imperfect information comes from the fact that strategies beh and bfg are absent from the model. Indeed, the events $[e] \cup [h]$ and $[f] \cup [g]$ are not conditions for player 1 in that model. But we can consider other possibilities, too, as we illustrate next.

The approach described above does not only make a conceptual difference, it can also have practical advantages in terms of modeling. To illustrate, consider again the game in the example, and suppose we want to model a situation where player 2 is uncertain about whether player 1 has perfect or imperfect information about her (player 2’s) move. This,

of course, requires adding strategies *beh* and *bfg* to the model. The latter is accomplished, for instance, by adding two more states, with respective paths *e* and *h*, forming a fourth type of player 1, and two more states, with respective paths *f* and *g*, forming a fifth type of player 1.¹⁷ In this case, we could still say that the second type of player 1 has imperfect information, if conditional on choosing neither *a* nor *beg*, she assigns probability one to her third type, who chooses *bfg*. Accordingly, to capture the idea that an imperfect information type cannot resort to strategies unavailable to her, we can weaken the definition of rationality, by comparing the actual strategy only with those that have positive conditional probability, given the event that the strategy is not the actual one (instead of comparing it to each other reduced strategy in the model). That is, under this weaker definition, an imperfect information type behaves rationally if her unconditional expected payoff is no less than her conditional expected payoffs, given each strategy chosen by other types who also have imperfect information.

6 Related models

6.1 Conditioning as the result of learning and updating

Battigalli and Siniscalchi (1999) studied conditional probability in a product type space à la Harsanyi (1967-68). Each agent's type is associated with a family of probability measures over basic states and the types of the *other* agents. Formally, in the model there is a set S of *external states*, a family $\mathcal{H}$ of nonempty subsets of S called *relevant hypotheses*, and a set of types T_i for each agent i . For each type t_i of agent i , there is a family $(\mu_i(t_i)(\cdot | H))_{H \in \mathcal{H}}$ of probability measures over $S \times T_{-i}$ satisfying conditional normality and the chain rule (when each hypothesis is viewed as a subset of $S \times T_{-i}$ in the obvious way). The authors note, that since in their model an agent's beliefs about her own type are not formalized, their analysis is based on the implicit assumption that the agent is certain of her own type for any given hypothesis. That is, for every agent i , type t_i and hypothesis H , the measure $\mu_i(t_i)(\cdot | H)$ is implicitly viewed as a measure on $S \times T_i \times T_{-i}$ that puts probability one on $S \times \{t_i\} \times T_{-i}$.

Formalizing the belief of each agent about his own type allows a formal statement of the said assumption and a comparison of their model with ours. For simplicity, assume that there are only two agents, 1 and 2, and that S , T_1 and T_2 are finite sets. Let $\Omega = S \times T_1 \times T_2$ and

¹⁷In the standard approach this would require considerably (and artificially) enlarging the game. In particular, we would need an initial move by Nature, choosing whether the actual game will be one with perfect or imperfect information. In the enlarged game, a strategy for player 1 would specify her action in case she has imperfect information, and also her action (as a function of her information) if she has perfect information. By contrast, the approach we propose here is simpler and perhaps more direct and explicit, as it does not require changing the *game* as we consider alternative informational assumptions. These assumptions, whether they concern observability of moves or subjective beliefs of any other kind, can be all accommodated by changing the informational model, that is, the *types*.

let $\mathcal{F}$ denote the product algebra of events on Ω . Let $\mathcal{C}$ be the family of events of the form $H \times T_1 \times T_2$, where $H \in \mathcal{H}$. Finally, for each agent i , assume a family of operators $B_i^p(\cdot | \cdot)$ mapping each event and condition into an event of the form $S \times E_i \times T_{-i}$, where $E_i \subseteq T_i$. By a straightforward modification of our proofs, we can verify that, for each agent i , the family $(\mu_i(t_i)(\cdot | H))_{H \in \mathcal{H}}$ corresponds to the family of operators $B_i^p(\cdot | \cdot)$ if and only if the latter satisfy axioms (P1)–(P5), (PN), (PC), and the axiom $B_i^p(E | C) \subseteq B_i^1(B_i^p(E | C) | D)$ for all $p \in [0, 1]$, $E \in \mathcal{F}$, and $C, D \in \mathcal{C}$. We may refer to this axiom as *strong introspection*: the agent is certain that her beliefs are the same, for every hypothesis.

The use of conditions in Battigalli and Siniscalchi (1999) and in this paper are diametrically opposed. Not only do we allow conditioning on events in the belief field, we require that all such events be conditions. Moreover, the axiom of echo implies that, conversely, all conditions are epistemic events. Thus, we capture the idea that conditional statements mean conditioning on the agent knowing the condition (see Corollary 3). In Battigalli and Siniscalchi (1999), conditions are assumed to be events that concern external, non-epistemic states. In effect, even without assuming it, this would follow as a result of strong introspection. More precisely, suppose that $B_i^1(E)$ was allowed as a relevant hypothesis. Then, $B_i^p(\neg E) \subseteq B_i^1(B_i^p(\neg E) | B_i^1(E))$ would be an instance of the axiom. But for $p > 0$, the righthand side of this inclusion is the empty set, and therefore $B_i^p(\neg E)$ is also empty, which implies that $B_i^1(E)$ is the whole space. That is, the conditional belief in the righthand side of the inclusion is, in fact, the unconditional belief. Thus, strong introspection prevents conditioning on nontrivial events about one's beliefs. The restriction of relevant hypothesis to non-epistemic events is inherent to the model.

The contradiction between echo and strong introspection corresponds to a basic difference in the interpretation of conditioning. This becomes particularly evident in the game-theoretic application of the two models.¹⁸ A model for a game, that satisfies strong introspection, is best understood as a model of belief *updating*. Conditions are information sets of the game, unconditional belief is belief at the initial node, and a conditional belief is the updated belief the player *would* hold at a (reached or unreached) information set, should she *learn* that it has been reached. These conditional beliefs are unrelated to actual beliefs held when the information set is actually reached. A player who is initially certain to choose an action, when conditioning on another action, does not imagine being a type that she is not, that is, a

¹⁸See also Battigalli and Siniscalchi (2002) and Battigalli, Di Tillio, and Samet (2013). In the latter paper, unlike in the former and like in the game-theoretic application of our work, strategies are not taken as primitive objects, but are instead constructed from players' conditional beliefs, while states of the world only specify the actual play. However, this is the only difference with respect to the cited papers of Battigalli and Siniscalchi. In particular, Battigalli, Di Tillio, and Samet (2013) still assume a product type space, external conditions, and (implicitly) strong introspection. The paper applies the model of Battigalli and Siniscalchi to the case where the external events concern play paths rather than strategies, its main question being whether strategies can be defined, and the existing analysis somehow replicated, based on play paths and conditional beliefs. While this is also one of the purposes of our discussion of the game-theoretic application of our model, the two exercises are based on underlying models that, as we argue in this section, turn out to be quite different.

type who is certain to choose the second action.¹⁹ By contrast, in the game-theoretic application of our model, conditioning on an unchosen action (a set of reduced strategies, hence an epistemic event under planning) means looking at counterfactual types who do choose that action, rather than looking at the belief that the same type of the player would hold, should she find herself at the information set that follows that action. While both models are static, in that beliefs are not indexed by time, ours also has a truly static interpretation. In particular, it has no explicit or implicit, formal or informal interpretation in terms of learning or updating. Indeed, the echo axiom would be inappropriate for this interpretation.

6.2 Hypothetical knowledge

A non-probabilistic version of epistemic conditioning is studied in Samet (1996). Conditional knowledge is described in that paper by a *hypothetical knowledge* operator on a state space that associates with each pair of events $H \neq \emptyset$ and E the event $K^H(E)$. To ease the comparison to our paper we denote $K^H(E)$ by $K(E | H)$ and the unconditional knowledge $K(\cdot | \Omega)$ by $K(\cdot)$. In what follows we compare the conditional knowledge operator K to the conditional certainty operator B^1 in this paper.

Seven axioms, (K1*)-(K7*), characterize a structure of the state space in Samet (1996). Except for the truth axiom (K7*), $K(E) \subseteq E$, they all either correspond to special cases of our axioms on conditional belief or follow from these axioms. In particular, axiom (K1*), $K(E | H) = K(K(E | H))$, which reflects introspection, is an instance of the equality in Proposition 3. Axioms (K2*) and (K3*) are $K(E | H) = K(K(E) | H)$ and $\neg K(E | H) = K(\neg K(E) | H)$. The first axiom corresponds to the instance of (Det) for $p = 1$. The second, follows from (Det).²⁰

The structure defined by these axioms has two elements. The first is a partition Π of the state space. The second is a *hypothesis transformation function* τ which assigns to each $\pi \in \Pi$ and hypothesis H an element of Π , $\pi' = \tau(\pi, H)$, such that $\pi' \cap H \neq \emptyset$, and $\pi = \pi'$ whenever $\pi \cap H \neq \emptyset$. The conditional $K(E | H)$ is true at π , that is, $\pi \subseteq K(E | H)$, when $K(E)$ is true at $\tau(\pi, H)$. The partition Π turns out to be a partition into types. Thus, in all the states in an element $\pi \in \Pi$ the conditional knowledge is the same. Thus, we may refer to π as a type, just as we do in the case of beliefs.

Because of axioms (K2*) and (K3*), which correspond to (Det), the structure of the type space in Samet (1996) has similar features to the one studied here under (Det). Consider the restriction of K to the epistemic field, namely the field generated by the partition Π . In

¹⁹Let E be the event that player i chooses a certain action. Then the following is an instance of the strong introspection axiom: $B_i^1(\neg E) \subseteq B_i^1(B_i^1(\neg E) | E)$. That is, after choosing an action she was initially certain not to choose, the player does not “forget” her initial beliefs.

²⁰Indeed, (Det) implies its negative version, $\neg B^p(E | C) \subseteq B^1(\neg B^p(E) | C)$. To see the latter, note that $\omega \notin B^p(E | C)$ implies $\omega \in B^q(\neg E | C)$ for some $q > 1 - p$. Hence $\omega \in B^1(B^q(\neg E) | C)$ by (Det), and therefore $\omega \in B^1(\neg B^p(E) | C)$ because $B^q(\neg E) \subseteq \neg B^p(E)$.

this case, $K(E | H)$ is true in π if $\tau(\pi, H) \subseteq H$, which follows from the requirement that $\tau(\pi, H) \cap H \neq \emptyset$, and $\tau(\pi, H) \subseteq E$. Thus, the events known *given* H for some given type π are those events that are known *unconditionally* for the type $\tau(\pi, H)$ in H .

Compare this to a type space as defined in this paper that satisfies (Det), and consider the restriction of B^1 to the epistemic field. In this case, certainties given C for π are the unconditional certainties for the most probable type, according to the type π , in C . This condition is similar to the one described in the previous paragraph, in that the conditional epistemic attitude is *determined* by the unconditional attitude of a *single* type.²¹

However, the truth conditions for B^1 and K differ in that for the first, the single type in the condition is determined by some order on types, while in the second it is selected arbitrarily by the function τ . This is due to the fact that axioms (PN) and (PC), which guarantee the hierarchy of types, have no counterpart in Samet (1996). Moreover, these two axioms are the reason why (Det) implies that the field of conditions is the epistemic field. In Samet (1996), this is no longer true, as the field of conditions is the whole power set.

6.3 Counterfactual conditionals

Type spaces for conditional probability can be considered as an extension of the models of counterfactual conditionals suggested by Lewis (1973). To see this, we consider a type space that satisfies (Int) and for which the condition field and the belief field coincide, that is, $\mathcal{E} = \mathcal{C}$. These two assumptions imply, by Corollary 1, that for each π the hierarchy $(S_1^\pi, \dots, S_{k^\pi}^\pi)$ satisfies $S_1^\pi = \pi$.

Consider now the restriction of the operator $B^1(\cdot | \cdot)$, which is defined on $\mathcal{F} \times \mathcal{E}$, to $\mathcal{E} \times \mathcal{E}$. With this restriction both the domain and range of B^1 are measurable with respect to the belief field $\mathcal{E}$ which is generated by Π . Thus, we can view B^1 as an operator on the state space with elements that are the members of Π . In this context we refer to these elements as *epistemic states*. The condition for $\pi \in \Pi$ to be in $B^1(E | C)$ is that $t^\pi(E | C) = 1$. By Claim 1, this is equivalent to $t^\pi(E | C^+) = 1$, where C^+ is the t^ω -positive part of C . Thus, π is in $B^1(E | C)$ when $S_{i_C}^\pi \subseteq E$, where i_C is the smallest index i for which $S_i^\pi \subseteq C$.

We now describe the structure delineated in the previous paragraph using the terminology of counterfactual conditionals. Rather than writing $B^1(E | C)$, we write $C \hookrightarrow E$, with the intended reading “if C then E ”. We think of the hierarchy $(S_1^\pi, \dots, S_{k^\pi}^\pi)$ as a partial order of the epistemic states, expressing closeness to π . Thus, the first element in the hierarchy, which is π , is the closest to π and the types in S_i^π are closer to π than those in S_k^π with $k > i$. We call a union of the form $\cup_{j=1}^i S_j^\pi$, a *sphere*. The family of spheres centered at π is denoted $\mathcal{S}^\pi$. Using this terminology, the truth condition for the conditional $C \hookrightarrow E$, described in the previous paragraph, is as follows. The conditional holds true for the epistemic state π (that

²¹Halpern (1999a) considers relaxations of (K3*) which makes τ a correspondence rather than a function, which implies that unconditional knowledge is not determined by a single type.

is π is in $C \hookrightarrow E$) when E holds in all the epistemic states in the intersection of C with the smallest sphere in $\mathcal{S}^\pi$ that intersects C non-vacuously. The description of the sphere system model, and the truth condition for the counterfactual conditional operator $\hookrightarrow$, are those given in Lewis (1973).

When the type space satisfies (Det) then each hierarchy consists of single types, or, using the term adopted in this subsection, of single epistemic states. The hierarchy at π is a simple ordering of epistemic states with π being the first. In this case π is in $C \hookrightarrow E$ when E contains the closest epistemic state to π , in the ordering associated with π . This model was proposed for counterfactual conditionals by Stalnaker (1968).

Conditional probability can be viewed as an extension of counterfactual conditionals. It provides us with a family of conditional operators that can be denoted by $\hookrightarrow_p$, where $C \hookrightarrow_p E$ is $B^p(E | C)$. We have shown that the restriction of the operator $\hookrightarrow_1$, denoted above as $\hookrightarrow$, to epistemic states is a counterfactual conditional. The axioms of (Echo) and (Det) extend the principle of truth condition of $\hookrightarrow_1$ to the family of probabilistic conditional operators $\hookrightarrow_p$ as follows. Whether the conditional $C \hookrightarrow_1 E$ is true in some state is answered by asking whether E is true, where C is used to select the states at which we check the truth of E . These are the states in C that are closest to the given state or, in the terminology of conditional probability, the most probable states in C . Analogously, whether the probabilistic conditional $C \hookrightarrow_p E$ is true in some state, is answered by asking whether a probabilistic statement about E is true, where C is used to select the states at which we check the truth of the statement, in the same manner that these states are selected for $\hookrightarrow_1$.

7 Final remarks

7.1 Introspection, echo, or determination?

The requirement that players know or are certain of their own beliefs or knowledge, as expressed in the axiom of introspection, is universally accepted in game theory and economic theory. The axioms of echo and determination both express the idea that the conditional beliefs are related to the unconditional beliefs held at the conditioning event. This idea is new, and its very formalization is made possible only because in a conditional belief type space, unlike in a standard belief type space, conditional probabilities rather than just probabilities are associated with each state.

Determination is stronger than echo, and it requires players to be very precise about their conditional beliefs. For example, consider a player with three possible choices, a, b and c, who chooses a. We ask the player: “Suppose you had not chosen a, would you have chosen b or c?” If determination holds, then the only possible answer is either b or c. If just echo holds, then the player’s answer can be that she is not sure, for example, that b and c could be

equally likely.²² Thus, echo gives more latitude in modelling, by allowing a wider range of reasonable answers to hypothetical questions.

7.2 The infinite case

Some of our analysis extends directly to the case of an infinite state space Ω , provided it is separable, that is, assuming a countably generated family of events $\mathcal{F}$. For example, in that case Theorem 1 and Proposition 2 remain true as stated—in particular, $\Pi(\omega)$ is still an event. In effect, these conclusions follow immediately from the analysis in Samet (2000), in particular from Theorem 3 in that paper. We also conjecture that analogues of our results concerning echo and determination, in particular Theorem 2 and Theorem 3, hold under appropriate technical conditions.

Besides presenting possible technical difficulties, the infinite case also raises a conceptual issue related to the assumption that the agent conditions on *events*. Indeed, with an infinite state space one could find it more appropriate to condition on a family of σ -fields that are not necessarily of the form $\{C, \neg C\}$, that is, not necessarily bipartitions—of course, this has to do with the very notion of conditional probability proposed by Rényi (1955), not directly with our use of that notion. Such *compatible families of conditional probabilities* (i.e., ones obeying the chain rule) have been studied, for example, by Sokal (1981).

7.3 Multiple agents

The extension of our analysis to the case of several agents is immediate—in effect, we have already appealed to it, when discussing the game-theoretic example in Section 5. Thus, with agents $1, \dots, N$, a conditional belief type space is a tuple $(\Omega, \mathcal{F}, \mathcal{C}_1, \dots, \mathcal{C}_N, t_1, \dots, t_N)$ where $\mathcal{C}_i \subseteq \mathcal{F}$ and $t_i : \Omega \rightarrow \Delta(\Omega, \mathcal{F}, \mathcal{C}_i)$ are, respectively, the condition field and type function of agent i , and $t_i(\cdot)(E | C)$ is assumed to be measurable with respect to $\mathcal{C}_i$ for each $E \in \mathcal{F}$ and $C \in \mathcal{C}_i$. Clearly, all of our results apply verbatim to each player separately.

Some additional facts can be also easily established. For instance, it can be easily seen that if echo holds for each player, and an event C is a condition for each player, that is, it belongs to $\mathcal{C}_i$ for every player i , then for each player conditioning on C is the same as conditioning on C being common knowledge (see Footnote 5). That is, the elements of the meet of the agents' partitions are the only events on which all agents can condition. This, in turn, gives an agreement theorem: at no state ω is it possible that all agents have the same conditional beliefs, given the condition C that there is common knowledge of disagreement on the probability of some event. Indeed, Corollary 5 implies that under echo, $t_i(\omega)(\cdot | C)$ is a prior of the unconditional type space for each agent i , hence a common prior.

²²This remark is similar in spirit to our earlier observation (Section 5.1) that in a game, echo alone pins down a player's reduced strategy (assuming planning) but not necessarily a strategy, in the sense that behavior at nodes excluded by the reduced strategy itself is not deterministically specified.

Finally, we remark that extending the analysis to the case of an infinite state space Ω is particularly important in the multi-agent setting, where constructing a universal belief space as in [Mertens and Zamir \(1985\)](#) can be a concern. The main challenge with this task lies in the fact that in considering higher and higher orders of belief, both the family of events *and* (due to echo) the family of conditions become larger. This issue is not present in the hierarchical construction of [Battigalli and Siniscalchi \(1999\)](#), because in that construction the family of conditions is exogenously given (see our discussion in [Section 6.1](#)).

Appendix

Proof of Theorem 1

It is easy to check that if a family of operators corresponds to a type function then it must satisfy (P1)–(P5), (PN), and (PC). Indeed, (P1)–(P5) follow from property (A) of conditional probability measures and the fact that these range in $[0, 1]$, whereas (PN) and (PC) follow from (N) and (C), respectively. To show the converse, we need the following preliminary result.

Lemma 1. *Let $(B^p)_{p \in [0,1]}$ be a family of operators satisfying (P1)–(P4) and (PN). Fix $C \in \mathcal{C}^+$. For every $p \in [0, 1]$ and $E, F \in \mathcal{F}$ such that $E \subseteq F$, $B^p(E | C) \subseteq B^p(F | C)$. Moreover, for every $E \in \mathcal{F}$ and $p, r \in [0, 1]$ such that $r > p$, $B^r(E | C) \subseteq B^p(E | C)$.*

Proof. The first claim follows from (P2), setting $q = 0$ and using (P1). From the first claim and (PN) it follows that $B^1(\Omega | C) = \Omega$. Thus, setting $p = 1$ and $E = \Omega$ in (P4), we have $B^q(\emptyset | C) = \emptyset$ for all $q \in (0, 1]$. Letting $F = \Omega$ and $q = r - p$ in (P3), the second claim follows. ■

Fix a family of operators $(B^p)_{p \in [0,1]}$ satisfying (P1)–(P4), (PN), and (PC). We now define a function t on Ω that assigns to each $\omega \in \Omega$ a function $t^\omega(\cdot | \cdot) : \mathcal{F} \times \mathcal{C}^+ \rightarrow [0, 1]$. Then we show that t is a type function. Finally, we prove that t is the unique type function to which $(B^p)_{p \in [0,1]}$ corresponds. For every $\omega \in \Omega$, $E \in \mathcal{F}$, and $C \in \mathcal{C}^+$ define $I(\omega, E, C) = \{p \in [0, 1] : \omega \in B^p(E | C)\}$. By (P1), this set is nonempty, as 0 belongs to it. By (P5), it has a maximum. Thus, for each $\omega \in \Omega$ define $t^\omega(\cdot | \cdot)$ by letting $t^\omega(E | C) = \max I(\omega, E, C)$ for every $E \in \mathcal{F}$ and $C \in \mathcal{C}^+$. To prove that t is a type function, we now fix $\omega \in \Omega$ and prove that t^ω satisfies (N), (A), and (C).

For every $C \in \mathcal{C}^+$, since $1 \in I(\omega, C, C)$ by (PN), we have $t^\omega(C | C) = 1$. Thus, t^ω satisfies (N). To prove that it satisfies (A), fix any $C \in \mathcal{C}^+$ and $A, B \in \mathcal{F}$ such that $A \cap B = \emptyset$. Let $p = t^\omega(A | C)$ and $q = t^\omega(B | C)$. Then, by the first claim in Lemma 1, $\omega \in B^p(A | C) \cap B^q(B | C) \subseteq B^p(A | C) \cap B^q(\neg A | C)$. Since the righthand side is nonempty, we conclude by (P4) that $p + q \leq 1$. By letting $E = A \cup B$ and $F = A$ in (P2), we obtain $B^p(A | C) \cap B^q(B | C) \subseteq B^{p+q}(A \cup B | C)$. As the lefthand side contains ω , it follows that

$t^\omega(A \cup B | C) \geq p + q$. Thus, $t^\omega(\cdot | C)$ is superadditive. Moreover, if $p + q = 1$ we must have $t^\omega(A \cup B | C) = p + q$. Therefore, it suffices to prove subadditivity for the case $p + q < 1$. Fix any $p' > p$ and $q' > q$ with $p' + q' \leq 1$. Then $\omega \in \neg B^{p'}(A | C) \cap \neg B^{q'}(B | C)$ and hence, by (P3), $t^\omega(A \cup B | C) < p' + q'$. As this is true for all such p', q' , we conclude that $t^\omega(A \cup B | C) \leq p + q$. This concludes the proof that t^ω satisfies (A). To prove that t^ω satisfies (C), fix any $E \in \mathcal{F}$ and $C, D \in \mathcal{C}^+$ with $E \subseteq D \subseteq C$. Let $p = t^\omega(E | D)$ and $q = t^\omega(D | C)$. Then $\omega \in B^p(E | D) \cap B^q(D | C)$ and hence, by (PC), $\omega \in B^{pq}(E | C)$. Thus, $t^\omega(E | C) \geq pq = t^\omega(E | D)t^\omega(D | C)$. This also holds if we replace E by $\neg E \cap D$, so $t^\omega(\neg E \cap D | C) \geq t^\omega(\neg E \cap D | D)t^\omega(D | C)$. Neither of these inequalities can be strict, because by adding them we would obtain, by the normality of $t^\omega(\cdot | D)$ and the additivity of $t^\omega(\cdot | C)$ and $t^\omega(\cdot | D)$, the contradiction $t^\omega(D | C) > t^\omega(D | D)t^\omega(D | C) = t^\omega(D | C)$. Thus, $t^\omega(E | C) = t^\omega(E | D)t^\omega(D | C)$. This shows that t^ω satisfies (C). The proof that t is a type function is complete.

To see that $(B^p)_{p \in [0,1]}$ corresponds to t , note that for all $p \in [0, 1]$, $E \in \mathcal{F}$, and $C \in \mathcal{C}^+$, we have $B^p(E | C) \subseteq \{\omega \in \Omega : t^\omega(E | C) \geq p\}$ by the definition of t , while the opposite inclusion holds by the second claim in Lemma 1. To establish uniqueness, let $\tilde{t}$ be a type function, and suppose that $(B^p)_{p \in [0,1]}$ corresponds to $\tilde{t}$. Fix $E \in \mathcal{F}$ and $C \in \mathcal{C}^+$, and let $p = \tilde{t}^\omega(E | C)$. Then $\omega \in B^p(E | C)$, and hence $p \in I(\omega, E, C)$. But for $q > p$, we have $\tilde{t}^\omega(E | C) < q$ and hence $q > \max I(\omega, E, C)$. Thus, $\tilde{t}^\omega(E | C) = \max I(\omega, E, C) = t^\omega(E | C)$.

Proof of Proposition 1

The following two lemmas imply Proposition 1, by the finiteness of Ω .

Lemma 2. *For each state ω , let $\mathcal{B}_\omega$ be the family of all $B \in \mathcal{B}$ such that $\omega \in B$. Then, $\Pi(\omega) = \cap_{B \in \mathcal{B}_\omega} B$.*

Proof. The inclusion of $\Pi(\omega)$ in $\cap_{B \in \mathcal{B}_\omega} B$ follows immediately from (1) and the definition of Π . For the opposite inclusion, fix a state $\omega' \in \cap_{B \in \mathcal{B}_\omega} B$. Since the type function t induces the family (B^p) , it follows that for all $p \in [0, 1]$, $E \in \mathcal{F}$, and $C \in \mathcal{C}^+$, if $t^\omega(E | C) \geq p$, then $t^{\omega'}(E | C) \geq p$. By normality and additivity, the probability measures $t^\omega(\cdot | C)$ and $t^{\omega'}(\cdot | C)$ must coincide. Thus, $\omega' \in \Pi(\omega)$. ■

Lemma 3. *For each $B \in \mathcal{B}$, $B = \cup_{\omega \in B} \Pi(\omega)$.*

Proof. For all $\omega \in \Omega$, $E \in \mathcal{F}$, $C \in \mathcal{C}^+$ and $p \in [0, 1]$, if $\omega \in B^p(E | C)$ then $\Pi(\omega) \subseteq B^p(E | C)$, by definition of Π . Thus, each $B \in \mathcal{B}$ is the union of the elements of Π that are contained in B . ■

References

- Aumann, R. J. (1976), “Agreeing to disagree,” *Annals of Statistics*, 4, 1236–1239.
- Battigalli, P., A. Di Tillio, and D. Samet (2013), “Strategies and interactive beliefs in dynamic games,” in *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress*, Econometric Society Monographs.
- Battigalli, P. and M. Siniscalchi (1999), “Hierarchies of conditional beliefs and interactive epistemology in dynamic games,” *Journal of Economic Theory*, 88, 188–230.
- Battigalli, P. and M. Siniscalchi (2002), “Strong belief and forward induction reasoning,” *Journal of Economic Theory*, 106, 356–391.
- Billingsley, P. (1995), *Probability and Measure*, John Wiley & Sons.
- Blume, L., A. Brandenburger, and E. Dekel (1991), “Lexicographic probabilities and choice under uncertainty,” *Econometrica*, 59, 61–79.
- de Finetti, B. (1936), “Les probabilités nulles,” *Bulletins des Science Mathématiques (première partie)*, 60, 275–288.
- Fagin, R. and J. Y. Halpern (1994), “Reasoning about knowledge and probability,” *Journal of the ACM*, 41, 340–367. A preliminary version appears in *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, 1988.
- Fagin, R., J. Y. Halpern, and N. Megiddo (1990), “A logic for reasoning about probabilities,” *Information and Computation*, 87, 78–128.
- Feinberg, Y. (2005), “Subjective reasoning—dynamic games,” *Games and Economic Behavior*, 52, 54–93.
- Gaifman, H. (1986), “A theory of higher order probabilities,” in *Theoretical Aspects of Reasoning about Knowledge: Proc. 1986 Conference*, pp. 275–292, Morgan Kaufmann.
- Halpern, J. Y. (1999a), “Hypothetical knowledge and counterfactual reasoning,” *International Journal of Game Theory*, 28, 315–330.
- Halpern, J. Y. (1999b), “Set-theoretic completeness for epistemic and conditional logic,” *Annals of Mathematics and Artificial Intelligence*, 26, 1–27.
- Halpern, J. Y. (2010), “Lexicographic probability, conditional probability, and nonstandard probability,” *Games and Economic Behavior*, 68, 155–179.
- Halpern, J. Y., D. Samet, and E. Segev (2009), “Defining knowledge in terms of belief: The modal logic perspective,” *The Review of Symbolic Logic*, 2, 469–487.

- Hammond, P. J. (1994), "Elementary non-Archimedean representations of probability for decision theory and games," in *Patrick Suppes: Scientific Philosopher; Volume I*, edited by P. Humphreys, pp. 25–49, Kluwer, Dordrecht, Netherlands.
- Harsanyi, J. C. (1967-68), "Games with incomplete information played by Bayesian players, I–III," *Management Science*, 14, 159–182, 320–334, 486–502.
- Heifetz, A. and P. Mongin (2001), "Probability logic for type spaces," *Games and Economic Behavior*, 35, 31–53.
- Keynes, J. M. (1921), *A Treatise on Probability*, Macmillan, London.
- Lewis, D. (1973), *Counterfactuals*, Basic Blackwell Ltd.
- Meier, M. (2012), "An infinitary probability logic for type spaces," *Israel Journal of Mathematics*, 192, 1–58.
- Mertens, J.-F. and S. Zamir (1985), "Formulation of Bayesian analysis for games with incomplete information," *International Journal of Game Theory*, 14, 1–29.
- Miller, D. (1966), "A paradox of information," *British Journal for the Philosophy of Science*, 17, 59–61.
- Monderer, D., D. Samet, and L. S. Shapley (1992), "Weighted values and the core," *International Journal of Game Theory*, 21, 27–39.
- Myerson, R. B. (1986a), "Axiomatic foundations of bayesian decision theory," Center for Mathematical Studies in Economics and Management Science, Northwestern University, Discussion Paper No. 671.
- Myerson, R. B. (1986b), "Multistage games with communication," *Econometrica*, 54, 323–358.
- Popper, K. R. (1934), *Logik der Forschung*, Julius Springer Verlag, Vienna.
- Popper, K. R. (1968), *The Logic of Scientific Discovery*, Hutchison, London, 2nd edition. The first version of this book appeared as *Logik der Forschung*, 1934.
- Rényi, A. (1955), "On a new axiomatic theory of probability," *Acta Mathematica Hungarica*, 6, 285–335.
- Rényi, A. (1956), "On conditional probability spaces generated by a dimensionally ordered set of measures," *Theory of Probability and Its Applications*, 1, 55–64.
- Samet, D. (1996), "Hypothetical knowledge and games with perfect information," *Games and Economic Behavior*, 17, 230–251.

- Samet, D. (1998a), “Common priors and separation of convex sets,” *Games and Economic Behavior*, 24, 172–174.
- Samet, D. (1998b), “Iterated expectations and common priors,” *Games and Economic Behavior*, 24, 131–141.
- Samet, D. (1999), “Bayesianism without learning,” *Research in Economics*, 53, 227–242.
- Samet, D. (2000), “Quantified beliefs and believed quantities,” *Journal of Economic Theory*, 95, 169–185.
- Sokal, A. D. (1981), “Existence of compatible families of proper regular conditional probabilities,” *Probability Theory and Related Fields*, 56, 537–548.
- Stalnaker, R. (1968), “A theory of conditionals,” in *Studies in Logical Theory*, edited by N. Rescher, pp. 98–112, Oxford: Blackwell.