



HAL
open science

Communication, common knowledge and consensus

Lucie Ménager

► **To cite this version:**

Lucie Ménager. Communication, common knowledge and consensus. Economics and Finance. Université Panthéon-Sorbonne - Paris I, 2006. English. NNT: . tel-00174147

HAL Id: tel-00174147

<https://theses.hal.science/tel-00174147>

Submitted on 21 Sep 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ PARIS I PANTHÉON-SORBONNE

Année 2006

N° attribué par la bibliothèque

2006 P006 PA01 0036

Thèse pour le doctorat de Sciences Economiques

soutenue publiquement par

Lucie Ménager

le 15 septembre 2006

COMMUNICATION, CONNAISSANCE COMMUNE ET CONSENSUS

Directeur de thèse

Monsieur Jean-Marc Tallon - Directeur de recherche au CNRS

Jury

Monsieur David Encaoua - Professeur à l'Université Paris 1

Madame Françoise Forges - Professeur à l'Université Paris-Dauphine

Monsieur Frédéric Koessler - Chargé de recherche au CNRS

Monsieur Nicolas Vieille - Professeur à l'Ecole HEC

L'université Paris I Panthéon-Sorbonne n'entend donner aucune approbation, ni improbation aux opinions émises dans cette thèse; elles doivent être considérées comme propres à leur auteur.

Remerciements

L'une des premières choses que m'a dite mon directeur est qu'une thèse se fait "avec les autres". Quatre ans après, je me rends compte à quel point c'est vrai, et je profite de ces pages pour remercier tous ceux qui m'ont guidée et accompagnée pendant ces quatre années.

Jean-Marc Tallon fut tout simplement un parfait directeur de thèse. C'est un euphémisme de dire qu'il fut à l'écoute, attentif, réactif, efficace et impliqué. Cette thèse a énormément bénéficié de ses relectures nombreuses et exigeantes. Surtout, il a la qualité rare de prendre au sérieux sans jamais mépriser, et de considérer que les idées n'ont pas moins de valeur lorsqu'elles viennent d'un doctorant. En bref, Jean-Marc a le don d'orienter sans éteindre la petite flamme de confiance qu'on a en soi, bien souvent vacillante au cours d'une thèse.

Je remercie vivement les membres de mon jury. Ils ont fait de nombreux commentaires sur mon travail à l'occasion d'une pré-soutenance, que j'ai essayé de mettre à profit de mon mieux. Françoise Forges et Nicolas Vieille ont accepté d'être les rapporteurs de ma thèse. C'est un grand honneur pour moi et je leur en suis très reconnaissante. En m'accueillant au sein de l'Ecole Doctorale, David Encaoua avait pris le temps de discuter avec moi de mon sujet de thèse. Je le remercie d'avoir accepté de faire partie de mon jury, et que la boucle soit ainsi bouclée. Enfin, tout au long de ma thèse, j'ai eu de nombreuses discussions avec

Frédéric Koessler, qui m'a accordé du temps sans compter. Je le remercie aujourd'hui pour son aide et son attention.

Pendant ces quatre années de thèse, j'ai fait partie du groupe de travail Choix, Information et Décision, véritable petite famille dont les membres m'ont soutenue, critiquée, découragée puis re-encouragée. Un proverbe chinois dit "Suis les conseils de celui qui te fait pleurer, pas de celui qui te fait rire". Notre groupe de travail est assurément un lieu où l'on suit beaucoup les conseils... Si Jean-Marc y fut symboliquement mon père de thèse, Jean-Christophe Vergnaud y fut mon oncle, dans le genre oncle grognon. Tel un oracle vivant au fond d'une grotte aux abords hostiles, il donne des réponses mystérieuses à ceux qui ont osé l'approcher en quête de connaissance. Michèle Cohen y tenait le rôle de grand-mère, et ses nombreux descendants, Mégléna Jéléva, Johanna Etner, Thibault Gajdos et Raphaël Giraud ont continué à être attentifs aux plus jeunes. Je remercie particulièrement Raphaël Giraud, Eric Danan, Jérôme Mathis, Olivier Baguelin et Olivier Tercieux. Nicolas Houy a été mon véritable camarade de thèse, partenaire de travail et de 400 coups. Sans lui tout aurait été différent. Enfin, ils sont le rappel constant que quatre années se sont écoulées...la Relève du groupe de travail! Merci à Morgane Tanvé, Benoît Menoni, Julie Moschion, Jeanne Hagenbach et Marie-Pierre Darnies pour leur présence attentive et sérieuse.

Bien qu'il soit extrêmement peu probable qu'il lise un jour ces lignes et qu'il les comprenne, il me faut remercier John Geanakoplos qui m'a remonté le moral dans un moment de doute. Cela ne lui coûtait pas grand chose et il ne se rappelle probablement pas de moi, mais ses phrases d'encouragement ont été décisives.

Le mois de juillet était caniculaire, le mois d'août était désert, et pourtant quelques

bonnes âmes ont bien voulu relire certains chapitres de ma thèse. Merci beaucoup à Olivier B. et Olivier T., à Christine, qui a emmené mon Chapitre 1 au Portugal, et à Yvon Hémerly, qui a eu la gentillesse de corriger quelques fautes d'anglais (j'en ai sûrement laissé).

Je remercie évidemment Tonia pour son efficacité au quotidien, et Elda pour sa prise en charge rassurante des tracasseries administratives de la fin de thèse. Je remercie également Joël Faron, véritable petit lutin de la MSE, toujours présent, toujours efficace, et surtout toujours très gentil. Jo, n'oublie pas tes gants!

Je veux également remercier ceux qui ont contribué indirectement à l'aboutissement de cette thèse, en participant à son commencement. Alain Chateaufort fut mon premier professeur de mathématiques en DEUG. Sa présence dans le pôle Théorie de la Décision avait quelque chose de très rassurant au début de ma thèse, alors que tout était si nouveau. Bernard Cornet et Jean-Marc Bonnisseau m'ont encadrée lors d'un stage qui a confirmé mon envie de faire de la recherche. Enfin, j'aurais pu, dans un autre monde, faire une thèse en Théorie de l'Equilibre Général. C'est indirectement grâce à Pierre Cahuc et à Antoine d'Autume que je me suis retrouvée en thèse au laboratoire EUREQua, et je les en remercie profondément. Car s'il est crucial de faire partie d'un groupe de travail exigeant (cf le proverbe chinois), il est aussi très important de discuter de sujets plus légers autour de la machine à café. A ce titre, ce fut un plaisir de côtoyer pendant ces quatre années Jean-Philippe Tropéano, Guillaume Hollard, Jean-Olivier Hairault, Antoine Bommier, Jean-Pierre Dugeon, Bertrand Wigniolle et Katheline Schubert.

Enfin, je veux remercier chaleureusement mes collègues doctorants. Bien sûr, on est un peu entassé, on étouffe (souvent), on s'énerve (parfois), mais surtout on se serre les coudes, on s'encourage, et on se reconforte. Le soutien et les attentions de ses camarades

doctorants n'ont pas de prix. Je pense d'abord aux "anciens", Djamel Stili, Hippolyte d'Albis et Lise Patureau, qui m'ont accueillie et ne m'ont pas oubliée une fois partis. Je remercie en particulier ma marraine de thèse Marie-Laure Michaud, qui m'a fait un trou dans la jungle des bureaux doctorants, et qui m'a enseigné quelques bonnes ficelles. Elle est ensuite devenue une amie, dont l'expérience et le recul ont été très précieux tout au long de ma thèse. Ensuite il y a les supers copines. Solenne la malicieuse, Audrey la touchante, précieuse complice des derniers mois, Céline la douce, pointe de fer dans un chausson de velours, et Andreea la cinglante, pied d'acier dans une botte de cuir. Puis les copains, Etienne, François, El Hadji, Falilou, Florent, Abuzer et Loïc. Il y a aussi les "p'tits jeunes", un peu turbulents mais tellement sympathiques. Morgane l'attentive, Sébastien l'attentionné, Julie l'attachante, Natacha et Marie-Pierre les rigolotes, Victor qui n'aime pas faire le point, Sumudu, Mohammed, Basak et Chanez qui sont toujours d'une si agréable compagnie. Enfin il y a Jeanne, arrivée tout droit d'Alsace dans un bec de cigogne, comme un cadeau pour ma dernière année. Il y a également tous les autres, trop nombreux pour les citer, qui ont fait de cette thèse une vraie aventure humaine.

Je les ai beaucoup négligés pendant la thèse, mais ils sont toujours là... Carole, Riad, Benoît, Elodie et Slimane sont des rocs d'amitié, depuis bientôt 10 ans sur les bancs de l'amphi K, et j'espère pour toujours. Merci à Christine, double casquette de super copine et co-auteur, à Gregory, dit binômet, et Cédric, dit Céd' La Balle : on s'est serré les coudes en DEA, et ça a finalement plutôt bien marché. Enfin il y a Jana, sous les bombes depuis le 12 juillet, et à qui j'ai pensé chaque jour de rédaction de cette thèse. Merci de m'avoir fait une fois de plus relativiser, cette fois ci sans rien dire.

Enfin, un grand merci aux membres de ma famille de m'avoir encouragée, chacun à sa manière, et de s'être intéressés à ce que je faisais. Merci à Xavier, à mes soeurs Charlotte,

Camille et Cécile, à leurs compagnons et à leurs enfants, à mes parents, à mes oncles et tantes, à mes cousins les crimous et mes cousines les vieilles chou, à mes grand-parents et particulièrement à ma chère grand-mère, qui a suivi avec attention les développements de ma thèse, et qui est certainement plus contente encore que moi qu'elle se termine enfin ("Quatre ans, c'est long quand même").

A mes soeurs

Note de synthèse

Introduction

Lorsqu'il fut interrogé sur les rapports établissant qu'il n'y avait pas de preuve d'un lien direct entre l'Iraq et certaines organisations terroristes, le secrétaire à la défense américain Donald Rumsfeld répondit:

Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know. And if one looks throughout the history of our country [...], it is the latter category that tend to be the difficult ones.

(Department of Defense News Briefing, 12 février 2002)¹

A travers ces quelques phrases apparemment sibyllines, Donald Rumsfeld exprimait l'idée suivante. Une bonne appréhension de sa connaissance et de sa méconnaissance est un

¹ *Je trouve toujours intéressants les rapports qui disent que quelque chose ne s'est pas produit car, comme on le sait, il y a des connus connus; il y a des choses dont on sait qu'on les connaît. On sait aussi qu'il y a des inconnus connus; c'est-à-dire qu'on sait qu'il y a certaines choses qu'on ne sait pas. Mais il y a aussi des inconnus inconnus, ceux dont on ne sait pas qu'on ne les connaît pas. Et lorsque l'on regarde l'histoire de notre pays [...], c'est cette dernière catégorie qui pose le plus de problèmes.*

atout majeur en stratégie militaire. Plus généralement, la *connaissance interactive*, c'est-à-dire la connaissance de la connaissance des autres, joue un rôle de première importance dans tout domaine impliquant des interactions stratégiques, en particulier en économie. Prenons l'exemple d'un duopole où deux entreprises se font concurrence en quantité. On suppose que si aucune entreprise ne connaît le niveau de production de l'autre, alors chacune produit la quantité qui maximise son espérance d'utilité. Appelons y_A et y_B ces quantités. Supposons maintenant que l'entreprise A ait l'opportunité d'observer secrètement la quantité produite par l'entreprise B , avant de choisir son niveau de production. A connaît la production de B , mais B ne le sait pas. B continue à produire y_B , tandis que A produit sa meilleure réponse face à y_B et augmente son profit. Imaginons maintenant que l'entreprise B apprenne que l'entreprise A a la possibilité d'observer son niveau de production avant de produire elle-même. L'entreprise B se retrouve alors dans la position d'un *leader* de Stackelberg : elle anticipe que l'entreprise A produira sa meilleure réponse face à sa production, et tire avantage de cette nouvelle connaissance. Ainsi, il n'est avantageux pour A de connaître le niveau de production de B que dans la mesure où B ne le sait pas. Ainsi, la stratégie optimale de B dépend de sa connaissance de y_A , et de sa connaissance de la connaissance de A sur y_B .

Un état remarquable de connaissance interactive est l'état de *connaissance commune*. On dit qu'un événement est connaissance commune dans un groupe d'individus lorsque chaque individu sait cet événement, que chacun sait que chacun sait cet événement, que chacun sait que chacun sait que chacun sait cet événement, etc. En économie, deux types d'enjeux sont liés au phénomène de connaissance commune. Le premier concerne les problèmes de coordination. Un problème de coordination est une situation caractérisée par le fait que les agents ont intérêt à participer à une action collective seulement si les autres

agents y participent aussi. Un moyen possible pour les individus de se coordonner pourrait être de communiquer un message du type “Participons tous à cette action collective”. Mais comme chaque personne ne veut participer que dans la mesure où les autres participent aussi, pour que le message permette aux agents de se coordonner, il faut non seulement que tous les individus soient au courant du message, mais aussi que tous sachent que tous sont au courant du message, etc. En un mot, il faut que le message soit connaissance commune dans le groupe. Considérons l'exemple suivant. Deux collègues décident d'aller boire un verre ensemble après leur journée de travail. Avant de monter dans le bus, ils se mettent d'accord pour descendre à l'arrêt n , pensant tous les deux que le bar se trouve à cet arrêt. Ils montent ensemble dans le bus, mais se trouvent séparés par le grand nombre de voyageurs. Le collègue A se retrouve complètement en tête du bus, et le collègue B complètement en queue. Lorsque les portes s'ouvrent à l'arrêt $n - 1$, les deux collègues se rendent compte que l'arrêt le plus proche du bar est en fait celui-là. Que vont-ils faire ? Descendre tout de suite alors qu'ils s'étaient mis d'accord pour descendre à l'arrêt d'après, et risquer ainsi de se retrouver seul ? Supposons qu'au moment où B réalise l'erreur, il cherche A des yeux mais n'arrive pas à voir si A a remarqué qu'il faudrait descendre tout de suite. B sait que l'arrêt $n - 1$ est le bon arrêt, mais ne sait pas si A le sait aussi. Comme B veut avant tout passer du temps avec A , il décide de rester dans le bus. Supposons maintenant qu'au moment où il réalise l'erreur, B voit A en train de le chercher des yeux sans succès. B sait alors que le bon arrêt est l'arrêt $n - 1$, sait que A le sait aussi, mais ne sait pas si A sait qu'il sait. Il décide donc de rester dans le bus. Si, en revanche, les regards de A et B se croisent au moment où ils réalisent leur erreur, alors ils descendent tous les deux du bus. Lorsque A et B ont réalisé l'erreur, que B sait que A a réalisé l'erreur, mais ne sait pas si A sait qu'il a réalisé l'erreur, il n'est pas connaissance commune entre A et B

que le bon arrêt est en fait l'arrêt $n - 1$. Ainsi, A et B restent dans le bus. Le fait d'établir un contact visuel rend connaissance commune entre A et B le fait qu'ils ont remarqué leur erreur, et les conduit alors à descendre à l'arrêt $n - 1$.

Le deuxième enjeu économique de la connaissance commune concerne le phénomène de *consensus*. Dans un article fondateur, Aumann [1976] montra que des individus rationnels ne peuvent pas "s'accorder sur un désaccord" à propos de leurs croyances *a posteriori*, lorsqu'ils ont les mêmes croyances *a priori* (les croyances des agents étant formalisées par leurs probabilités sur les états du monde). Plus précisément, le *théorème d'accord* d'Aumann peut s'énoncer de la manière suivante. Supposons que deux agents aient les mêmes probabilités *a priori*. Si les probabilités *a posteriori* qu'ils attribuent à un événement donné sont connaissance commune entre eux, alors elles sont égales. L'importance de ce résultat tient au fait qu'il suggère que les asymétries d'information ont moins de pouvoir explicatif qu'on ne pourrait le penser à première vue. En effet, ce résultat montre qu'on ne peut pas expliquer des différences de croyances par des différences d'information privée, lorsque ces différences de croyances sont connaissance commune. En particulier, il implique que les échanges spéculatifs ne peuvent pas être expliqués uniquement par de l'asymétrie d'information. En effet, supposons que deux spéculateurs aient la même croyance *a priori*, et que tous les deux reçoivent une information privée contradictoire sur l'évolution future du cours d'une certaine action a . L'un croit que le cours de l'action va baisser, et souhaite par conséquent vendre de l'action a . L'autre croit que le cours de a va monter, et donc souhaite acheter de l'action a . Supposons que ces deux spéculateurs se rencontrent, et se mettent d'accord sur un prix d'échange. On peut alors penser que l'échange est connaissance commune entre eux, c'est-à-dire qu'il est connaissance commune entre eux que l'un accepte d'acheter, et l'autre accepte de vendre. Cela signifie qu'il est connaissance com-

mune entre eux que les deux spéculateurs ont des croyances opposées sur l'évolution future du cours de l'action a . Or cela n'est pas possible, d'après le résultat d'Aumann. Ainsi, pour expliquer les échanges spéculatifs, il faut soit supposer que les agents sont imparfaitement rationnels, soit supposer qu'ils ont des croyances *a priori* différentes.

Le résultat d'Aumann a donné lieu à une vaste littérature de la part d'économistes, de philosophes et de logiciens, portant aussi bien sur les fondements épistémiques de ce type de réflexions que sur des extensions du résultat d'Aumann dans différentes directions. Cette thèse présente trois contributions originales à cette littérature, que nous avons choisi d'appeler la littérature *Agreeing to Disagree*, en référence au résultat d'Aumann. Elle se compose de six chapitres. Les deux premiers chapitres présentent le cadre formel et la littérature. Les trois chapitres suivants sont des contributions originales. Enfin, le dernier chapitre compare les propriétés des règles de décision utilisées dans les chapitres 3 et 4 avec celles des règles existant dans la littérature.

Chapitre 1 : Modéliser la connaissance

Introduction

La littérature *Agreeing to Disagree* se base sur une modélisation particulière de la connaissance et de la connaissance commune, particulièrement appropriée dans les cadres interactifs, où l'on a besoin de modéliser la connaissance qu'ont les agents de la connaissance des autres agents. Le premier chapitre de cette thèse est ainsi consacré à une présentation détaillée du modèle de connaissance utilisé dans la littérature et dans cette thèse. Ce modèle, appelé *structure à la Aumann*, consiste en un ensemble d'*états du monde* et en une *partition d'information* pour chaque agent. Un état du monde est une description complète

du monde, incluant les faits objectifs comme “Il pleut” ainsi que les faits de connaissance comme “Je sais qu’il pleut”. Une partition d’information est une partition de l’ensemble des états du monde, représentant la manière dont les agents sont informés sur le monde en chacun des états. Dans le chapitre 1, on insiste sur l’interprétation controversée de la notion d’état du monde, et on apporte des éléments de réponse à certains pseudo-paradoxes liés à ces problèmes d’interprétation. Parmi les questions posées par l’interprétation des structures à la Aumann, deux peuvent particulièrement remettre en cause l’interprétation des résultats de notre thèse. Les partitions individuelles sont-elles “connaissance commune” entre les agents? Les états du monde évoluent-ils lorsque les agents révisent leur information privée? Pour clarifier ces questions, on se propose de revenir aux fondements logiques des structures à la Aumann, c’est-à-dire au modèle de connaissance utilisé de longue date en logique épistémique.

Les structures à la Aumann

On considère un groupe d’individus raisonnant à propos d’un monde qui peut être décrit en termes de *faits objectifs*, tels que “Il pleut”, et de *faits de connaissance*, tels que “Je sais qu’il pleut”. Un *état du monde* consiste en la liste de tous les faits (objectifs et de connaissance) qui sont vrais dans ce monde. A la manière de la théorie des probabilités, l’approche utilisée en économie pour modéliser la connaissance se base sur les *événements*, qui sont des ensembles d’états du monde. Plus précisément, un événement est le *champ* d’une propriété, c’est-à-dire l’ensemble des états du monde dans lesquels cette propriété est vraie, et on identifie les événements avec les propriétés dont ces événements sont le champ. Par exemple, l’ensemble des états du monde dans lesquels la propriété “Il pleut” est vraie est identifié à l’événement “Il pleut”. On note Ω l’ensemble des états du monde et

2^Ω l'ensemble des événements de Ω . L'ensemble vide correspond à une propriété qui n'est jamais vraie, c'est-à-dire à une contradiction, et l'ensemble Ω à une propriété toujours vraie, c'est-à-dire à une tautologie. L'inclusion entre des événements correspond à une implication matérielle entre propriétés. Enfin, on note $\neg E$ la négation de l'événement E , *i.e.* $\neg E := \Omega \setminus E$.

Une structure à la Aumann consiste en un ensemble d'états du monde Ω et une partition d'information Π_i pour chaque individu i .² Une partition d'information est une fonction $\Pi : \Omega \rightarrow 2^\Omega$ telle que pour tout $\omega \in \Omega$:

$$(i) \quad \omega \in \Pi(\omega)$$

$$(ii) \quad \omega' \in \Pi(\omega) \Rightarrow \Pi(\omega') = \Pi(\omega)$$

La partition d'information d'un agent représente sa relation de possibilité, c'est-à-dire les états que l'agent croit possibles dans chaque état du monde. Plus précisément, $\Pi_i(\omega)$ est l'ensemble des états jugés possibles par l'agent i lorsque l'état ω se réalise. L'information objective, correspondant à la réalisation de l'état du monde ω , apparaît déformée à l'agent i par sa partition d'information Π_i .

Modéliser la connaissance dans les structures à la Aumann.

Il y a deux manières équivalentes de représenter la connaissance dans une structure à la Aumann. La première est de lister l'ensemble des événements dont l'individu sait qu'ils se sont réalisés, étant donnée l'information qu'il a à sa disposition. On dit qu'un individu muni d'une partition Π_i sait l'événement E sous l'état ω si et seulement si

$$\Pi_i(\omega) \subseteq E$$

²Les structures à la Aumann sont des structures d'information particulières où les *correspondances de possibilité* satisfont trois axiomes, de réflexivité, transitivité et euclidianité. Voir le Chapitre 1 de la thèse.

En d'autres termes, on dit qu'un individu connaît un événement si cet événement est réalisé en chacun des états que l'individu croit possibles.

La deuxième manière consiste à énumérer, pour chaque événement, l'ensemble des états du monde dans lesquels l'individu sait l'événement. Pour cela, on définit un opérateur de connaissance individuel $K : 2^\Omega \rightarrow 2^\Omega$. Etant donné un individu i et un événement E , $K_i(E)$ est l'événement " i sait E ". La partition d'un agent peut être construite à partir de son opérateur de connaissance, de même que l'opérateur de connaissance peut être construit à partir de la partition de l'agent, *via* la relation suivante :

$$K(E) = \{\omega \in \Omega \mid \Pi(\omega) \subseteq E\}$$

On note $\neg K$ l'opérateur de "méconnaissance". Ainsi, $\neg K(E)$ est l'événement "l'agent ne sait pas E ". Comme $K(E)$ est un événement particulier, on peut lui appliquer l'opérateur de méconnaissance $\neg K$. $\neg K(K(E))$ est alors l'événement "l'agent ne sait pas qu'il sait E ". Cet événement est-il réalisable, c'est-à-dire existe-t-il un état ω tel que $\omega \in \neg K(K(E))$? Plus généralement, quelles propriétés de la connaissance suppose-t-on implicitement lorsqu'on considère des structures à la Aumann? On montre dans le chapitre 1 que lorsque la connaissance des agents est représentée par une partition de l'ensemble des états du monde, alors leur opérateur de connaissance vérifie les cinq propriétés axiomatiques suivantes :

A1 Axiome de conscience : les individus sont capables d'identifier l'ensemble des mondes possibles. Ainsi, ils ne sont jamais surpris par une contingence qu'ils n'auraient pas anticipée.

A2 Omniscience logique : lorsqu'un individu sait un événement, il sait aussi toutes les implications logiques de cet événement, c'est-à-dire tous les événements contenant

cet événement.

A3 Axiome de vérité : bien qu'il soit possible que des agents ne connaissent pas tous les événements réalisés, lorsqu'ils savent un événement, alors cet événement est réalisé.

Autrement dit, les individus ne peuvent pas se tromper, en sachant des choses fausses.

A4 Introspection positive : un individu ne peut pas savoir un événement sans savoir qu'il le sait. Autrement dit, lorsqu'un agent sait un événement, alors il sait qu'il le sait.

A5 Introspection négative : un individu ne peut pas ignorer un événement sans savoir qu'il l'ignore. Autrement dit, lorsqu'un individu ne sait pas un événement, il sait qu'il ne le sait pas.

Dans le chapitre 1, on présente en détail et on discute la force et la plausibilité de ces cinq propriétés axiomatiques.

Modéliser la connaissance commune dans les structures à la Aumann

On attribue la définition informelle de la connaissance commune au philosophe Lewis [1969]. Un événement E est dit connaissance commune dans un groupe d'individus si chacun sait E , sait de plus que chacun sait E , sait que chacun le sait, et ainsi de suite. Dans un article fondateur à plus d'un titre, Aumann [1976] fut le premier à formaliser la notion de connaissance commune. Définissons d'abord l'*union*³ d'un ensemble de partitions.

Definition 1 (Union de partitions) *L'union des partitions Π_1, \dots, Π_n , notée M , est la plus fine grossissement commun de ces partitions, c'est-à-dire la plus fine partition telle que*

$$\forall i, \forall \omega, \Pi_i(\omega) \subseteq M(\omega).$$

³Traduit de l'anglais *meet*.

Aumann [1976] montre qu'un événement est connaissance commune en un état ω s'il contient la cellule de la partition M qui contient ω :

Proposition 1 (Aumann [1976]) *Considérons un groupe d'agents \mathcal{N} , chaque agent $i \in \mathcal{N}$ étant muni d'une partition Π_i , et soit M l'union des partitions $(\Pi_i)_{i \in \mathcal{N}}$. Un événement E est connaissance commune en ω si et seulement si $M(\omega) \subseteq E$.*

Aumann montra ainsi que l'union d'un groupe de partitions individuelles est la partition de connaissance commune dans le groupe. Selon cette définition, il ne suffit que d'un nombre fini d'étapes pour vérifier si un événement est connaissance commune, lorsque le nombre d'agents et les partitions d'information sont finis. On montre dans le chapitre 1 que la connaissance commune satisfait les mêmes propriétés que la connaissance individuelles.

Controverses

Dans la section 3 du chapitre 1, on présente deux controverses qui furent à la base de critiques violentes à l'égard de la littérature *Agreeing to Disagree*. La première porte sur ce que les individus savent de l'information des autres, et se décline en deux questions.

1. Est-ce que les partitions des individus sont "connaissance commune" entre eux?
2. Est-il nécessaire que les partitions individuelles soient "connaissance commune" pour que l'union des partitions soit effectivement la partition de connaissance commune?

A l'aide de la formalisation de la connaissance utilisée de longue date par les logiciens et les philosophes, on montre que la première question n'a pas de sens telle qu'elle est posée. En effet, les individus ne peuvent pas "connaître" les partitions des autres : la connaissance des individus porte sur des faits, objectifs ou de connaissance, et les partitions sont une représentation de cette connaissance. Ainsi, la connaissance des individus

ne peut pas porter sur un outil de représentation, nécessairement artificiel, et essentiellement utile au modélisateur. A l'aide des mêmes outils (Lismont et Mongin [1994]), on explique que l'union des partitions individuelles est également une représentation sémantique de l'opérateur de connaissance commune, et qu'aucune hypothèse supplémentaire n'est nécessaire pour utiliser cette représentation.

La deuxième controverse est liée à l'utilisation dynamique qui est faite des structures à la Aumann dans la littérature *Agreeing to Disagree*. En effet, on suppose dans cette littérature que les agents révisent leur information selon la règle suivante. Soit Π^t la partition d'information d'un agent à la date t , et $m : \Omega \times \mathbb{N} \rightarrow \mathcal{M}$ une fonction de message, où $m(\omega, t)$ désigne le message entendu par l'agent à la date t et sous l'état ω . L'ensemble des états jugés possibles par l'agent à la date $t + 1$ sous l'état ω est l'ensemble des états que l'agent pensait possibles à la date t , qui n'ont pas pu être éliminés sur la base du message $m(\omega, t)$. Formellement,

$$\Pi^{t+1}(\omega) = \Pi^t(\omega) \cap \{\omega' \in \Omega \mid m(\omega', t) = m(\omega, t)\}$$

L'agent n'a pas la même connaissance à la date $t + 1$ qu'à la date t . Or les états du monde décrivent entre autres choses l'état de connaissance des agents. La controverse prend ainsi la forme de la question suivante.

3. Peut-on considérer que les états du monde ne changent pas entre la date t et la date $t + 1$?

Pour que la règle de révision utilisée ait un sens, il faut qu'elle permette d'exprimer une connaissance nouvelle des mêmes états du monde. Pour lever ce pseudo-paradoxe, on se réfère aux travaux récents de Bonanno [2004] et Board [2004]. Ils montrent que dans un cadre dynamique, les états du monde doivent non seulement décrire l'état de connaissance

des agents à la date présente, mais aussi l'arrivée d'une information nouvelle, et les états de connaissance après réception de la nouvelle information. Dans ce cas, les états du monde ne changent pas, et les partitions d'information à la date $t + 1$ représentent effectivement une connaissance nouvelle d'un même incertain.

Chapitre 2 : Une synthèse de la littérature *Agreeing to Disagree*.

Introduction

Aumann [1976] montra que la connaissance commune dans un groupe d'agents peut être représentée par une partition particulière, construite à partir des partitions individuelles. Dans cet article fondateur, il prouva de plus la propriété suivante. Si deux agents ont les mêmes croyances *a priori*, alors leurs croyances *a posteriori* pour un événement donné ne peuvent pas être différentes, si elles sont connaissance commune entre eux. Ainsi, des différences d'opinion actées de tous, par exemple sur les perspectives des marchés financiers, ne peuvent pas être expliquées uniquement par le fait que les agents détiennent des informations privées différentes. Dans le chapitre 2, on présente le *théorème d'accord* d'Aumann, ainsi qu'une synthèse des travaux réalisés à la suite de ce résultat. En section 2.5, on présente les principales critiques adressées à l'encontre de cette littérature, notamment par Moses et Nachum [1990].

Le résultat fondateur : Agreeing to Disagree

Le résultat d'Aumann s'énonce de manière élégante en anglais, mais se traduit difficilement en français. Il dit en substance que "*Rational agents cannot agree to disagree*", c'est-à-dire que des agents rationnels ne peuvent pas s'accorder sur leurs désaccords.

Théorème 1 *Soient deux agents A et B munis de partitions Π_A et Π_B , et soit $E \subseteq \Omega$ un événement donné. Si A et B ont la même probabilité a priori P sur Ω , et s'il est connaissance commune en l'état ω que $P(E \mid \Pi_A(\omega)) = p_A$ et $P(E \mid \Pi_B(\omega)) = p_B$, alors $p_A = p_B$.*

L'hypothèse de probabilité *a priori* commune a été, et est encore, violemment critiquée en théorie des jeux (voir notamment Lipman [1995], Gul [1996]). Dans le chapitre 2, on souligne le fait que la contribution majeure du résultat d'Aumann ne nécessite pas l'hypothèse d'*a priori* commune. En effet, la preuve du théorème utilise l'argument suivant. Soit P_i la probabilité *a priori* de l'agent i , et considérons un événement quelconque E . Si $P_i(E \mid \Pi_i(\omega))$ est connaissance commune en ω , alors $M(\omega) \subseteq \{\omega' \in \Omega \mid P_i(E \mid \Pi_i(\omega')) = P_i(E \mid \Pi_i(\omega))\}$ par définition. Or par construction, $M(\omega)$ est une union, nécessairement disjointe, de cellules de Π_i . Par conséquent, $P_i(E \mid M(\omega)) = P_i(E \mid \Pi_i(\omega))$. Ainsi, si la probabilité *a posteriori* d'un agent est connaissance commune, alors cette probabilité aurait été la même, si l'agent avait conditionné sa probabilité non pas sur son information privée $\Pi_i(\omega)$, mais sur l'information publique $M(\omega)$. Cette propriété est souvent interprétée comme le fait que *la connaissance commune neutralise l'information privée*. En effet, supposons que les probabilités *a posteriori* des individus d'un groupe soient toutes connaissance commune dans ce groupe. On sait alors que pour tout individu i , $P_i(E \mid \Pi_i(\omega)) = P_i(E \mid M(\omega))$. Supposons maintenant que les membres du groupe "permutent" leur partitions d'information. La partition de connaissance commune ne change pas, par conséquent, les probabilités individuelles doivent encore être égales à $P_i(E \mid M(\omega))$. Ainsi, la connaissance commune des probabilités *a posteriori* individuelles implique que les probabilités des individus ne reflètent pas les différences d'information privée de chacun.

La littérature Agreeing to Disagree

Les travaux réalisés à la suite du résultat d'Aumann peuvent être classés en deux groupes, selon la question à laquelle ces travaux répondent. Les travaux du premier groupe généralisent au sens strict le résultat d'Aumann, selon deux directions. La première direction est de considérer que les agents expriment non pas des probabilités *a posteriori*, mais des décisions. La seconde direction est de considérer que la connaissance commune porte non pas sur les décisions individuelles, mais sur un agrégat, ou statistique, de ces décisions. Ainsi, les travaux généralisant le théorème d'Aumann répondent à la question suivante :

Question I : *Quelles conditions garantissent que la connaissance commune d'une statistique des décisions individuelles implique que les décisions des individus ne reflètent pas leur information privée?*

Les travaux du deuxième groupe tentent d'expliquer l'émergence de situations de connaissance commune, au travers de processus d'échange d'informations entre les agents. Plus précisément, ces travaux répondent à la question suivante :

Question II : *Quelles conditions garantissent que la communication des décisions individuelles conduit à la situation de connaissance commune de ces décisions?*

Concernant cette question, on choisit d'articuler les résultats de la littérature selon la nature publique ou privée de la communication. En effet, lorsque la communication est publique, la connaissance commune des décisions émerge en toute généralité. Lorsque la communication est non publique en revanche, la connaissance commune des décisions émerge *via* le consensus, qui est une situation où les agents prennent tous la même décision, et où cette décision partagée est connaissance commune entre eux. Par conséquent, il faut au moins supposer l'égalité des règles de décisions pour garantir l'émergence de la connaissance commune des décisions individuelles dans les protocoles non publics.

Chapitre 3 : Connaissance commune d'un agrégat de décisions.

Introduction

Le chapitre 3 est consacré à une étude de la question de type I. Aumann [1976] y apporta le premier un élément de réponse, en montrant que lorsque la connaissance commune porte sur les décisions individuelles, c'est-à-dire lorsque la statistique des décisions individuelles est la fonction identité, alors il suffit que les décisions des agents soient leurs probabilités *a posteriori* pour un événement. Dans le même cadre, Cave [1983] et Bacharach [1985] ont montré qu'il suffit que les règles de décision des agents satisfassent une condition plus générale dite de *stabilité par l'union*.⁴ Une fonction $\delta : 2^\Omega \rightarrow \mathcal{D}$ est stable par l'union si pour tous $E, E' \subseteq \Omega$ tels que $E \cap E' = \emptyset$, $\delta(E) = \delta(E') = d \Rightarrow \delta(E \cup E') = d$. Enfin, McKelvey et Page [1986] ont considéré le cas où la connaissance commune porte non pas sur les décisions individuelles, mais sur un agrégat, ou statistique, de ces décisions. Dans ce cadre, ils montrent qu'il suffit que 1) les individus aient tous la même croyance *a priori*, que 2) les décisions des agents soient leur probabilités *a posteriori* pour un événement donné, et que 3) la statistique des décisions individuelles soit *stochastiquement régulière*,⁵. A titre d'illustration, le résultat de McKelvey et Page implique que si les agents ont la même probabilité *a priori*, et si la moyenne des probabilités *a posteriori* qu'ils attribuent à un événement donné sont connaissance commune entre eux, alors ces probabilités sont égales.

Dans le chapitre 3, on considère le même cadre que celui de McKelvey et Page [1986], c'est-à-dire que la connaissance commune ne porte pas sur les décisions individuelles mais sur un agrégat de ces décisions, en supposant à l'instar de Cave [1983] et Bacharach [1985]

⁴Traduit de l'anglais *union consistency*.

⁵Traduit de l'anglais *stochastically regular*.

que les décisions des agents ne sont pas nécessairement leur probabilités *a posteriori*. On montre que si les règles de décision suivies par les agents satisfont une condition dite de *stabilité par l'union équilibrée*, et si la statistique des décisions individuelles satisfait une condition dite d'*exhaustivité*, alors la connaissance commune de la valeur de la statistique implique que les décisions des agents ne reflètent pas leur information privée. Les hypothèses de notre théorème nous permettent de considérer l'application suivante. Supposons que les membres d'une commission de recrutement se réunissent afin de décider d'engager quelqu'un. Supposons de plus que les membres de la commission aient tous comme objectif d'engager le candidat le plus qualifié, mais que chacun ait reçu une information privée sur les compétences des différents candidats. En arrivant à la réunion, chaque recruteur a ainsi sa propre opinion sur les capacités des candidats. Supposons qu'un sondage soit réalisé auprès de chacun des membres de la commission, et qu'on leur demande quel est selon eux le candidat le plus qualifié pour le poste. Nous montrons dans le chapitre 3 que si la valeur de ce sondage est connaissance commune, alors les membres de la commission doivent être tous d'accord sur le candidat à engager.

Le modèle

On considère un groupe d'agents \mathcal{N} , chaque agent i étant muni d'une partition Π_i de l'ensemble des états du monde Ω , supposé fini. La manière dont les agents prennent leur décisions est déterminée par une *règle de décision*, qui prescrit aux agents quelle décision prendre en fonction de leur information privée. Précisément, les agents suivent tous la même règle de décision $\delta : 2^\Omega \rightarrow \mathcal{D}$, où \mathcal{D} est l'ensemble de décision des agents. Comme Ω est fini, l'ensemble des situations d'information possibles l'est également, et on note $\{d_1, \dots, d_m\}$ cet ensemble de décisions possibles, avec $m < \infty$.

Les agents prennent leur décisions en fonction de leur information privée. Lorsque l'état ω se réalise, l'agent i est informé de $\Pi_i(\omega)$, et par conséquent prend la décision $\delta(\Pi_i(\omega))$. On note $\underline{\delta}_i(\omega)$ la décision prise par i en l'état ω , et $\underline{\delta}(\omega)$ le profil de décisions sous ω . Dans ce chapitre, on étudie l'effet que la connaissance commune d'une statistique Φ des décisions individuelles en un état donné ω aura sur le profil $\underline{\delta}(\omega)$. On dit que la statistique Φ est connaissance commune en ω si l'événement $\{\omega' \in \Omega \mid \Phi(\underline{\delta}(\omega')) = \Phi(\underline{\delta}(\omega))\}$ est connaissance commune en ω .

La condition que l'on impose sur la règle de décision suivie par les agents est la *stabilité par l'union équilibrée*. Avant de définir cette condition, définissons ce que l'on appelle une *famille équilibrée*, dont le sens diffère légèrement de celui de Shapley [1967].

Definition 2 Une famille non-vide $\mathcal{B} \subseteq 2^\Omega$ est équilibrée s'il existe une famille de coefficients positifs ou nuls $\{\lambda_S\}_{S \in \mathcal{B}}$, appelés coefficients d'équilibrage, tels que $\sum_{S \in \mathcal{B}} \lambda_S \mathbb{1}_{\omega \in S} = 1$ pour tout $\omega \in \bigcup_{S \in \mathcal{B}} S$.

Un exemple de famille équilibrée de $\Omega = \{1, 2, 3, 4, 5, 6\}$ est $\mathcal{B} = \{\{1, 2\}\{3, 4\}\{1, 2, 4\}\{1, 2, 3\}\}$, qui est équilibrée par rapport aux coefficients $\lambda_{1,2} = \lambda_{1,2,4} = \lambda_{1,2,3} = 1/3$ et $\lambda_{3,4} = 2/3$.

Nous pouvons maintenant donner la définition de la stabilité par l'union équilibrée.

Definition 3 Une règle de décision δ est stable par l'union équilibrée si et seulement si pour toute famille équilibrée \mathcal{B} , $\delta(S) = d \forall S \in \mathcal{B} \Rightarrow \delta(\bigcup_{S \in \mathcal{B}} S) = d$.

La stabilité par l'union équilibrée implique la stabilité par l'union de Cave [1983] et Bacharach [1985], qui impose que si un individu prend la même décision sachant deux événements disjoints, alors il prend encore la même décision sachant l'union de ces événements. Cependant, la stabilité par l'union équilibrée est satisfaite par des règles de décision

usuelles en économie, comme les règles *argmax*, dont une définition précise est donnée dans le chapitre 4.

La condition que l'on impose à la statistique des décisions individuelles est l'*exhaustivité*. On dit qu'une statistique Φ est exhaustive si elle est une transformation injective de la statistique Φ^* définie comme suit :

Definition 4 $\Phi^* : \{d_1, \dots, d_m\}^N \rightarrow \mathcal{N}^m$ est définie par $\Phi^*(x_1, \dots, x_N) = (\sum_{i=1}^N \mathbb{1}_{x_i=d_1}, \dots, \sum_{i=1}^N \mathbb{1}_{x_i=d_m})$.

En d'autres termes, une statistique exhaustive est une sorte de mesure de comptage des décisions individuelles. Typiquement, un sondage sur l'ensemble d'une population est une statistique exhaustive.

Le résultat

Dans le chapitre 3, on montre le résultat suivant. Considérons un groupe d'agents suivant tous la même règle de décision stable par l'union équilibrée. Si les agents du groupe ont connaissance commune d'une statistique exhaustive de leurs décisions individuelles, alors tous les agents doivent prendre la même décision.

Théorème 2 Soient δ une règle de décision stable par l'union équilibrée et Φ une statistique exhaustive. Pour tout $\omega \in \Omega$, si $\Phi(\underline{\delta}(\omega))$ est connaissance commune en ω , alors $\underline{\delta}_i(\omega) = \delta(M(\omega))$ pour tout i .

Dans le chapitre 3, nous expliquons les rôles joués par les conditions de stabilité par l'union équilibrée et d'exhaustivité dans l'établissement du résultat. On montre en particulier que ni la condition de stabilité par l'union de Cave [1983] ni celle de convexité de Parikh et Krasucki [1990] ne permettent de garantir le résultat pour plus de trois agents.

Enfin, on compare la condition d'exhaustivité avec la condition de régularité stochastique de McKelvey et Page.

Chapitre 4 : Consensus, communication et connaissance : une extension avec des agents Bayésiens

Introduction

Dans le chapitre 4, on présente une contribution originale aux travaux traitant de la question de type II. Geanakoplos et Polemarchakis [1982] furent les premiers à y apporter un élément de réponse. Ils montrèrent que si deux agents ont les mêmes probabilités *a priori*, et si ces agents révèlent et révisent leur probabilités *a posteriori* pour un événement donné, alors au bout d'un nombre fini d'étapes, ces agents finiront par avoir les mêmes probabilités *a posteriori*. Cave [1983] montra que des conditions suffisantes dans le cas d'un nombre quelconque d'agents, sont que le protocole de communication soit public et simultané, et que la règle de décision suivie par les agents soit stable par l'union. Parikh et Krasucki [1990] considèrent le cas plus général d'une communication éventuellement privée et séquentielle entre plusieurs agents. Ils supposent que les agents communiquent leurs décisions selon un protocole Pr , qui détermine les émetteurs et les récepteurs de la communication à chaque date. Ils montrent que deux conditions sont suffisantes pour que la communication conduise à un consensus : le protocole de communication doit être *équitable*,⁶ c'est-à-dire tel qu'aucun agent ne soit exclu de la communication, et la fonction dont les valeurs sont communiquées doit être *convexe*.⁷ Une fonction $\delta : 2^\Omega \rightarrow \mathcal{D}$ est convexe si pour tous $E, E' \subseteq \Omega$ tels que $E \cap E' = \emptyset$, il existe $\alpha \in]0, 1[$ tel que $\delta(E \cup E') =$

⁶Traduit de l'anglais *fair*.

⁷Traduit de l'anglais *convex*.

$\alpha\delta(E)+(1-\alpha)\delta(E')$. La condition de convexité est satisfaite notamment par les probabilités *a posteriori*. Parikh et Krasucki montrent que si la fonction est seulement *faiblement convexe*,⁸ alors la situation de consensus peut ne pas émerger dans certains protocoles équitables.

Il existe des espaces de décisions pour lesquels la condition de convexité de Parikh et Krasucki ne peut pas s'appliquer. Considérons un individu sur le point d'acheter une voiture. Son ensemble de décisions est {Acheter, Ne pas acheter}. Supposons qu'on renomme les décisions dans \mathbb{R} , avec 0 pour Acheter et 1 pour Ne pas acheter. La condition de convexité implique que s'il existe deux événements disjoints X et Y tels que $\delta(X) = 1$ et $\delta(Y) = 0$, alors $\delta(X \cup Y) \in]0, 1[$, ce qui ne correspond à aucune décision dans {Acheter, Ne pas Acheter}.

Dans le chapitre 4, on identifie une nouvelle condition sur la fonction dont les valeurs sont communiquées qui garantit que la communication conduise à un consensus dans tout protocole équitable. On montre que si les agents ont la même fonction d'utilité, et s'ils communiquent selon un protocole équitable l'action qui maximise leur espérance d'utilité, alors ils vont atteindre un consensus sur leur action optimale. Contrairement à la condition de convexité, notre condition, appelée *argmax*, s'applique à tout espace de décisions.

Même après avoir renommé leur image dans \mathbb{R} , les fonctions considérées dans ce chapitre ne sont pas toujours représentables par des fonctions faiblement convexes. De plus, il existe des fonctions convexes qui ne satisfont pas notre condition. Ainsi, la classe des fonctions *argmax* est d'intersection non vide avec la classe des fonctions faiblement convexes, mais il n'y a pas de relation d'inclusion entre elles. D'autre part, les fonctions *argmax* sont stables

⁸Une fonction $\delta : 2^\Omega \rightarrow \mathcal{D}$ est faiblement convexe si pour tout $E, E' \subseteq \Omega$ tels que $E \cap E' = \emptyset$, il existe $\alpha \in [0, 1]$ tel que $\delta(E \cup E') = \alpha\delta(E) + (1 - \alpha)\delta(E')$

par l'union quel que soit l'espace de décisions considéré.

Les hypothèses de notre théorème permettent de l'appliquer au problème de la diffusion d'innovations. Considérons un groupe de producteurs de blé en Beauce qui ont soudain l'opportunité de produire du maïs. Chacun s'est renseigné sur la convenance du sol beauceron pour l'exploitation du maïs, mais les fermiers ne peuvent pas partager leur information privée car ils travaillent toute la journée. Le seul contact qu'ils ont les uns avec les autres est un contact visuel : chacun est capable d'observer ses plus proches voisins, mais pas les voisins de ses voisins. En particulier, chaque fermier observe le choix de variété de ses voisins, mais pas celui des voisins de ses voisins, qui sont trop loin. Supposons que certains fermiers jugent préférable, au vu de leur information privée, de cultiver du maïs. Comment vont réagir les autres fermiers? Une situation où certains fermiers produiraient du blé et d'autres du maïs est-elle une situation d'équilibre stationnaire? Le théorème présenté dans le chapitre 4 montre qu'au bout d'un certain temps, tous les fermiers vont finir par produire la même variété de céréales, et que le choix de cette variété consensuelle sera connaissance commune entre eux.

Le modèle

On considère toujours un groupe fini d'agents \mathcal{N} , chaque agent i étant muni d'une partition d'information Π_i de l'ensemble des états du monde Ω , supposé fini. Les agents suivent une règle de décision $\delta : 2^\Omega \rightarrow \mathcal{D}$, et communiquent leurs décisions aux autres. On suppose que \mathcal{D} est un sous-ensemble compact d'un espace topologique quelconque. La manière dont les agents communiquent est définie par un protocole de communication, qui détermine les émetteurs et les récepteurs de la communication à chaque date. Plus précisément,

Definition 5 *Un protocole de communication Pr est une paire de fonctions $(s(\cdot), r(\cdot)) : \mathbb{N} \rightarrow 2^{\mathcal{N}} \times 2^{\mathcal{N}}$, telle que $s(t)$ et $r(t)$ sont respectivement les ensembles d'émetteurs et de récepteurs de la communication à la date t .*

On suppose que le protocole de communication est équitable.⁹ La question de type II se traduit alors par la suivante. Quelles conditions faut-il imposer sur le protocole de communication Pr et sur la règle de décision suivie par les agents pour arriver à un consensus, c'est-à-dire une situation où tous les agents prennent la même décision, et où cette décision partagée est connaissance commune entre eux?

On définit maintenant la condition argmax. Supposons que les agents aient tous la même fonction d'utilité $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$, dont la valeur dépend de l'action choisie $d \in \mathcal{D}$ et de l'état du monde réalisé. On suppose que $U(\cdot, \omega)$ est continue sur \mathcal{D} pour tout ω . Les agents communiquent l'action qui maximise leur espérance d'utilité, calculée par rapport à leur probabilité *a priori* P . Afin d'éviter les cas d'indifférence, on fait l'hypothèse que, pour tout événement, toutes les actions ont des utilités espérées différentes conditionnellement à cet événement.

[Hypothèse de non-indifférence] *Pour tout événement $F \subseteq \Omega$, $\forall d \neq d' \in \mathcal{D}$, $E(U(d, \cdot) | F) \neq E(U(d', \cdot) | F)$.*

Sans cette hypothèse, l'ensemble des actions optimales ne serait pas nécessairement un singleton, et il faudrait alors spécifier la manière dont les agents choisissent l'action communiquée dans cet ensemble. Cette hypothèse est cependant assez forte, et on discute des conséquences de son relâchement dans la conclusion du chapitre 4.

Definition 6 *Une fonction $\delta : 2^{\Omega} \rightarrow \mathcal{D}$ est du type argmax s'il existe une fonction*

⁹Voir la définition dans la thèse

d'utilité $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$ et une probabilité P sur Ω telles que pour tout $X \subseteq \Omega$,
 $\delta(X) = \operatorname{argmax}_{d \in \mathcal{D}} \mathbb{E}_P[U(d, \cdot) \mid X]$.

Supposons maintenant que Pr soit un protocole de communication donné. L'ensemble des états possibles pour l'agent i à la date t lorsque l'état du monde est ω est noté $C_i(\omega, t)$, et est défini par le processus récursif suivant :

$$C_i(\omega, 0) = \Pi_i(\omega)$$

$$C_i(\omega, t+1) = C_i(\omega, t) \cap \{\omega' \in \Omega \mid \delta(C_{s(t)}(\omega', t)) = \delta(C_{s(t)}(\omega, t))\} \text{ si } i = r(t),$$

$$C_i(\omega, t+1) = C_i(\omega, t) \text{ sinon.}$$

Le résultat

Dans le chapitre 4, on montre que si les agents suivent une règle de décision de type argmax , alors la communication des décisions individuelles selon un protocole équitable permet de créer un consensus de décisions. Autrement dit, on montre que pour tout ω , $\delta(C_i(\omega, t))$ admet une valeur limite qui ne dépend pas de i .

Théorème 3 *Il existe $T \in \mathbb{N}$ tel que pour tout ω , i , et tout $t, t' \geq T$, $C_i(\omega, t) = C_i(\omega, t')$. De plus, si le protocole est équitable, alors pour tout i, j , et pour tout ω , $\delta(C_i(\omega, T)) = \delta(C_j(\omega, T))$.*

On discute des propriétés des fonctions de type argmax dans le chapitre 6. Tout d'abord, une fonction argmax est clairement *stable par l'union* quel que soit l'espace de décision. Ensuite, une fonction argmax n'est pas forcément représentable par une fonction faiblement convexe, c'est-à-dire qu'étant donné f argmax , il peut ne pas exister une fonction bijective $g : \mathcal{D} \rightarrow \mathbb{R}$ telle que $g \circ f$ soit faiblement convexe. Si une telle fonction g existait, les propriétés de f et de $g \circ f$ en termes d'apprentissage et de consensus seraient les mêmes.

Ainsi, les fonctions argmax seraient des fonctions faiblement convexes particulières pour lesquelles le consensus émerge dans tout protocole équitable. On montre que ce n'est pas le cas dans le chapitre 6. Enfin, il existe des fonctions faiblement convexes qui ne peuvent pas être définies comme l' argmax d'une espérance d'utilité. Un tel exemple peut être trouvé dans Parikh and Krasucki [1990, p 185]: ils présentent une fonction f faiblement convexe telle qu'aucun consensus n'émerge dans certains protocoles équitables. On montre dans le chapitre 6 qu'il n'est pas possible de trouver une fonction d'utilité U et une probabilité P telle que cette fonction f soit l' argmax de l'espérance conditionnelle de U .

Chapitre 5 : Consensus, communication et ordre de parole : qui veut parler en premier?

Introduction

Issu d'un travail en collaboration avec Nicolas Houy, le chapitre 5 traite de la question de l'influence de l'ordre de parole sur la prise de décision en groupe. Considérons l'histoire suivante, inspirée de l'exemple bien connu des trois chapeaux. Alice et Bob sont assis l'un en face de l'autre, chacun portant un chapeau pouvant être rouge ou blanc. Supposons que les deux chapeaux soient rouges. Leur professeur annonce aux deux enfants qu'au moins un de leurs chapeaux est rouge, et leur demande s'ils connaissent la couleur de leur chapeau. Le fait qu'il soit maintenant connaissance commune entre les deux enfants qu'au moins un de leur chapeaux est rouge n'est pas suffisant pour leur permettre d'en inférer la couleur de leur propre chapeau. La seule manière pour eux de répondre à la question du professeur est de communiquer d'une manière ou d'une autre. Supposons qu'Alice dise à Bob qu'elle ne connaît pas la couleur de son chapeau. Bob comprend alors que son chapeau est rouge, puisque s'il avait été blanc, Alice aurait compris que son chapeau à elle était rouge. Si

maintenant Bob dit à Alice qu'il sait la couleur de son chapeau, Alice n'apprendra rien, car Bob l'aurait su que son chapeau eût été rouge ou blanc. En effet, Alice sait que si son chapeau était blanc, Bob aurait su que son chapeau à lui était rouge, d'après la remarque du professeur, et elle sait que si son chapeau était rouge, Bob aurait su que son chapeau à lui était rouge parce qu'elle ne savait pas la couleur de son chapeau. Si, en revanche, Bob dit en premier qu'il ne sait pas la couleur de son chapeau, alors Alice comprendra que son chapeau est rouge. Ainsi, si Alice veut apprendre la couleur de son chapeau, elle n'a pas intérêt à être la première à parler. Cette histoire illustre le fait suivant. A partir du moment où les agents communiquent de manière à apprendre de l'information, l'ordre de parole n'est pas anodin. Les processus de communications ne sont pas commutatifs, puisque différents ordres de parole conduisent à des issues différentes.

Parikh and Krasucki [1990] ont étudié le cas où les agents d'un groupe se communiquent la valeur d'une fonction f , selon un protocole de communication sur lequel ils se sont mis d'accord au préalable. Ils montrèrent que, si le protocole de communication est équitable, et si la fonction dont les valeurs sont communiquées est convexe, alors la communication conduit à un consensus sur la valeur de cette fonction. Le point de départ de ce travail fut l'observation que, dans le cadre de Parikh et Krasucki, différents protocoles ont des issues différentes, en termes de valeur du consensus atteint ainsi qu'en termes de montant d'information apprise par les agents au cours du processus. En particulier, il peut arriver qu'un agent apprennent plus d'information en communiquant avec les autres *via* un protocole α que *via* un protocole β . Il peut également arriver que les protocoles les plus informatifs ne soient pas les mêmes pour tous les agents. Ainsi, si les agents communiquent de manière à apprendre de l'information, ils peuvent être en désaccord quant au protocole de communication à utiliser.

La question étudiée dans le chapitre 5 est la suivante. On considère le même cadre que celui de Parikh et Krasucki, en faisant l'hypothèse supplémentaire que les agents communiquent afin d'apprendre de l'information. Implicitement, les agents sont des preneurs de décision, qui essaient d'être mieux informés au sens de Blackwell [1953] afin de prendre de meilleures décisions. Cette hypothèse implique que les agents ont des préférences dépendantes des états sur les protocoles de communication. Dans le chapitre 5, on étudie les inférences que peuvent faire des agents rationnels de la connaissance commune que certains d'entre eux sont en désaccord quant au protocole de communication à utiliser.

On observe que les deux situations suivantes sont possibles. Premièrement, il peut être connaissance commune dans un groupe d'agents que tous préfèrent le même ordre de parole. Deuxièmement, il peut être connaissance commune dans un groupe d'agents que certains d'entre eux sont en désaccord quant au protocole de communication à utiliser. Cependant, on montre le résultat surprenant que dans ce cas, la valeur de consensus de la fonction doit être la même, quel que soit le protocole utilisé.

Quel est le sens de ce résultat dans le cadre de l'exemple des producteurs de blé? On sait que, sous certaines conditions, les fermiers vont finir par tous produire du blé, ou tous produire du maïs. On sait aussi que la céréale "consensuelle", sur laquelle les fermiers vont tous tomber d'accord, dépend du protocole de communication utilisé, c'est-à-dire dans ce cadre de la manière dont les fermiers sont localisés les uns par rapport aux autres. Supposons que certaines fermes soient plus convoitées que d'autres, non parce que le terrain y est meilleur, mais parce que leur disposition permet d'apprendre beaucoup d'information (on peut penser par exemple à la ferme centrale, de laquelle on peut observer tous les autres fermiers). Supposons que deux fermiers veuillent acheter la même ferme, et que ce soit connaissance commune dans la région. Le résultat du Chapitre 5 implique que la

variété consensuelle sera la même, quel que soit le fermier qui possède la ferme.

Le modèle

On reprend (en le généralisant un peu) le cadre de Parikh et Krasucki, et on introduit les notations suivantes. Π_i^α désigne la partition d'information de l'agent i à l'équilibre du processus de communication, lorsque le protocole de communication est α . Π^α désigne la partition de connaissance commune à l'équilibre du processus, c'est-à-dire l'union des partitions individuelles Π_i^α , et $f(\Pi^\alpha(\omega))$ désigne la valeur consensuelle de f sous l'état ω , lorsque le protocole utilisé est α .

On sait que dans ce cadre, étant donné n'importe quel protocole α , sous les hypothèses d'équitabilité et de convexité, la communication des valeurs privées de la fonction f conduit à un consensus sur la valeur de f . La proposition suivante établit que cette valeur dépend du protocole de communication.

Proposition 2 *Il existe un modèle d'information $\langle \Omega, (\Pi_i)_i, f \rangle$ avec f convexe et deux protocoles équitables α, β pour lesquels il existe ω tel que $f(\Pi^\alpha(\omega)) \neq f(\Pi^\beta(\omega))$.*

Ce résultat peut être montré facilement pour des fonctions stables par l'union. Cependant, il n'avait pas été montré pour des probabilités *a posteriori*. Comme les probabilités *a posteriori* sont des fonctions stables par l'union, il eut pu être possible qu'il n'existât pas de modèle avec des probabilités *a posteriori* tel que l'ordre de parole compte. Dans le chapitre 5, on montre un exemple prouvant la proposition ??.

On fait l'hypothèse que les agents communiquent de manière à être mieux informés au sens de Blackwell [1953]. On sait qu'un partition Π est plus informative qu'une partition Π' si et seulement si Π est plus fine que Π' , c'est-à-dire si chaque cellule de Π est incluse

dans une cellule de Π' . Ainsi, on dit qu'un agent est mieux informé lorsqu'il communique avec un protocole α qu'avec un protocole β si sa partition d'information à l'équilibre du processus est plus fine avec α qu'avec β .

Cela induit pour chaque agent une préférence dépendante de l'état sur l'ensemble des protocoles. Avant que la communication n'ait lieu, l'ensemble des états pertinents pour l'agent i sous l'état ω est $\Pi_i(\omega)$. Ainsi, on dit qu'un agent préfère α à β sous l'état ω si Π_i^α est plus fine que Π_i^β en tout état que l'agent juge possible sous ω .

Definition 7 (Préférences) Soient α et β deux protocoles. L'ensemble des états du monde dans lesquels l'agent i préfère α à β est noté $B_i(\alpha, \beta)$, et est défini comme suit :

$$B_i(\alpha, \beta) = \{\omega \in \Omega \mid \forall \omega' \in \Pi_i(\omega), \Pi_i^\alpha(\omega') \subseteq \Pi_i^\beta(\omega') \text{ et } \exists \omega'' \in \Pi_i(\omega) \text{ s.t. } \Pi_i^\alpha(\omega'') \subset \Pi_i^\beta(\omega'')\}$$

Le résultat

On énonce maintenant le résultat principal du chapitre 5.

Théorème 4 Soit $\langle \Omega, (\Pi_i)_i, f \rangle$ un modèle d'information tel que f soit convexe, et soient α, β deux protocoles équitables tels que $\alpha \neq \beta$. Considérons $a_1, a_2, b_1, b_2 \in \{\alpha, \beta\}$, avec $a_1 \neq a_2$ et $b_1 \neq b_2$, et fixons $i \neq j$. Les assertions (1), (2) et (3) ne peuvent pas être vraies simultanément.

(1) $B_i(a_1, a_2)$ et $B_j(b_1, b_2)$ sont connaissance commune en ω .

(2) $\omega \in B_i(a_1, a_2) \cap B_j(b_1, b_2)$ et $a_1 = b_2$.

(3) $f(\Pi^\alpha(\omega)) \neq f(\Pi^\beta(\omega))$.

La signification de ce théorème dans l'exemple donné en introduction est la suivante.

- Si (1) et (2) sont vraies, c'est-à-dire s'il est connaissance commune en ω que Alice et Bob veulent tous les deux parler en premier (ou en deuxième), alors (3) est fausse, *i.e* la valeur de consensus de f en ω est la même, que Alice ou Bob parle en premier.

- Si (1) et (3) sont vraies, c'est-à-dire si les préférences d'Alice et de Bob concernant l'ordre de parole sont connaissance commune en ω , et si la valeur consensuelle de f dépend du fait qu'Alice ou Bob parle en premier, alors (2) est fausse, *i.e* Alice et Bob préfèrent le même protocole en ω .


- Si (2) et (3) sont vraies, c'est-à-dire si Alice et Bob veulent tous les deux parler en premier (ou en deuxième) en ω , et si la valeur consensuelle de f est différente selon que Alice ou Bob parle en premier, alors (1) est fausse, *i.e* les préférences d'Alice ou de Bob ne sont pas connaissance commune en ω .

Dans le chapitre 5, on montre que le résultat de ce théorème n'est pas dû au fait que deux des trois assertions ne peuvent pas être vraies simultanément. On montre aussi qu'il n'est pas dû au fait que lorsque (1) et (2) sont vraies, alors les partitions de connaissance commune à l'équilibre du processus sont les mêmes avec α et β .

Chapitre 6 : Une étude comparée de certaines propriétés des règles de décision

La relation entre la convexité, la convexité faible et la stabilité par l'union fut clarifiée par Parikh et Krasucki [1990] : l'ensemble des fonctions convexes est strictement inclus dans l'ensemble des fonction faiblement convexes, lui-même strictement inclus dans l'ensemble des fonctions stables par l'union. Dans le chapitre 6, on examine d'abord les liens entre les nouvelles conditions introduites dans la thèse, argmax et stabilité par l'union équilibrée, et les conditions de convexité et faible convexité. Ensuite, on précise si chacune des conditions garantit que 1) la connaissance commune d'une statistique exhaustive des décisions individuelles implique l'égalité des décisions, et que 2) la communication selon un protocole équitable conduit à un consensus.

On note **U** l'ensemble des fonctions stables par l'union, **BU** l'ensemble des fonctions stables par l'union équilibrée, **A** l'ensemble des fonctions argmax, **C** l'ensemble des fonctions convexes et **WC** l'ensemble des fonctions faiblement convexes. On montre les inclusions décrites par la figure ci-dessous.



On montre également que parmi les cinq conditions identifiées, la convexité et la condition argmax sont les seules garantissant que la communication conduise à un consensus dans tout protocole équitable, pour plus de trois agents. La stabilité par l'union équilibrée est la seule garantissant que la connaissance commune d'une statistique exhaustive des décisions individuelles implique un consensus, pour plus de trois agents.

Introduction

In a now famous response to questions about reports stating there was no evidence of a direct link between Iraq and terrorist organizations, US Defense Secretary Donald Rumsfeld said in 2002:

Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know. And if one looks throughout the history of our country [...], it is the latter category that tend to be the difficult ones.

(Department of Defense News Briefing, February 12 2002)

Donald Rumsfeld found an elegant way to express the idea that knowledge about knowledge is of primary importance in military intelligence. More generally, the notions of knowledge and knowledge about knowledge are fundamental in any area implying strategic interactions, in particular in economics.

Consider the case of two firms engaged in a duopolistic competition. Suppose that both duopolists do not know the quantity produced by the other. Each duopolist produces the quantity that maximizes his expected profit, say y_A and y_B . Suppose now that duopolist

A has the opportunity to secretly learn duopolist B 's output before choosing its own. A knows B 's output, but B does not know that A knows its output. B still produces y_B and A increases its profits by producing its best response to y_B . Suppose now that B learns that A will choose its output after observing its own output. B finds itself in the position of a Stackelberg leader. B anticipates A 's best response and takes advantage of this new knowledge. Therefore, A 's knowledge about B 's production is an advantage for A only if B does not know it.

This simple example shows that knowledge about others' knowledge, namely *interactive knowledge*, is at least as important as knowledge about objective facts such as production levels. A particular state of interactive knowledge is the state of *common knowledge*. Something is said to be common knowledge in a group of agents when everybody in the group knows it, everybody knows that everybody knows it, and so on, *ad infinitum*.

The notion of common knowledge

Common knowledge is a phenomenon which is inherent to much of social life. In order to communicate, people need to have some common knowledge of the language they use. In order to coordinate their behavior successfully, people require some common understanding. Everyday life is full of social *conventions*, such as driving on the right, defined by Lewis [1969] as strict coordination equilibria which agents follow on account of their common knowledge that they all prefer to follow these equilibria. Hume [1740] was perhaps the first to make explicit reference to the role of mutual knowledge in coordination. He argued that a necessary condition for coordinated activity was that agents all know what behavior to expect from one another. Much later, Littlewood [1953] presented some examples of common knowledge-type reasoning, and Schelling [1960] argued that something

like common knowledge was needed to explain certain inferences people make about each other. The first to give an explicit analysis of common knowledge was Lewis [1969], in his book *Convention*, and the first to provide a formal characterization of common knowledge was Aumann [1976] in his celebrated article *Agreeing to Disagree*. Assuming that individuals are informed about the world by an information partition, he showed that common knowledge can be represented by a particular partition, derived from individual partitions.

Let us give some motivating examples to illustrate a variety of ways in which the actions of agents depend crucially on their having, or lacking, certain common knowledge.

The mischievous father A father has two children, Alice and Bob. He gives to each of them an envelope, containing a certain amount of euro coins, and tells them that one envelope contains one more euro than the other, with Alice having an odd number of euro, and Bob an even number of euro. Each child observes the amount of money in his or her envelope, but not the amount in the other child's envelope. First, the father asks Alice and Bob to announce if he or she knows who has the larger amount of money. Of course, neither knows this, so none of them makes an announcement. Now, the father tells Alice and Bob that no envelope contains ten euro or more. Initially, nothing happens. But eventually one of the two announces that he or she is richer. Why? If Alice had 9 euro, she would immediately realize that Bob has 8 euro, and would therefore announce that she was the richest. But now if Bob has 8 euro and Alice does not announce that she is the richest, he will be able to infer that Alice must have 7 euro. By backward induction, one can infer that eventually, the richest child will realize that he or she is the richest.

The department store (Thomas Schelling, *The Strategy of Conflict*, p 54) When a man

loses his wife in a department store without any prior understanding on where to meet if they get separated, the chances are good that they will find each other. It is likely that each will think of some obvious place to meet, so obvious that each will be sure that it is “obvious” to both of them. One does not simply predict where the other will go, since the other will go where he predicts the first to go, which is wherever the first predicts the second to predict the first to go, and so *ad infinitum*. Not “What would I do if I were she?” but “What would I do if I were she wondering what she would do if she were I wondering what I would do if I were she...?” What is necessary is to coordinate predictions, to read the same message in the common situation, to identify the one course of action that their expectations of each other can converge on. They must “mutually recognize” some unique signal that coordinates their expectations of each other. We cannot be sure they will meet, nor would all couples read the same signal; but the chances are certainly a great deal better than if they pursued a random course of search.

Getting off the bus (adapted from Chwe [2001]) Say you and I are co-workers who ride the same bus home. Today the bus is completely packed and somehow we get separated, with you standing near the front door of the bus and me near the back door. Before we reach our usual stop, I notice a mutual acquaintance who yells from the sidewalk, “Hey you two! Come join me for a drink!” Joining this acquaintance would be nice, but we care mainly about each other’s company. The bus doors open; separated by the crowd, we must decide independently whether to get off. What should we do? Several situations are possible. Suppose that when our acquaintance yells out, I look for you but cannot find you: I am not sure whether you notice her or not and thus stay in the bus. We may both know that our acquaintance yelled but I do not know

that you know. Suppose now that when our acquaintance yells, I see you raise your head and look around for me, but I am not sure whether you manage to find me. I know about the invitation, I know that you know since I see you looking at me, but I stay in the bus because I do not know whether you know that I know about the invitation. Suppose finally that we do manage to make eye contact, we get off the bus.

Rebellion (from Jehl [1996]) For nearly thirty years, the price of a loaf of bread in Egypt was held constant. Anwar El-Sadat's attempt in 1977 to raise the price was met with major riots. Since then, one government tactic has been to quietly replace a fraction of the wheat flour with cheaper corn flour. Each person could notice that their own loaf tasted different, but be unsure about how many other people noticed. Changing the taste of the loaves is not the same public event as raising its price.

The three hats Imagine three pupils sitting in a circle, each wearing either a red hat or a white hat. Suppose that the three hats are red, and that the teacher tells the girls they are allowed to leave the room every time the school bell rings, if they know the color of their hat with confidence. No girl ever leaves the class room, since no girl can see her own hat. Suppose now that the teacher tells the girls that at least one hat is red, a fact which is well-known, since every girl can see two red hats in the room. The first two times the bell rings, no girl leaves the room. The third time the bell rings however, the three girls leave the room, knowing with confidence that they are all wearing a red hat. Why? By publicly announcing that at least one hat is red, the teacher makes the information common knowledge among the three girls. The fact that girl 3 does not leave the room the first time the ring bells informs the two

other girls that at least one of their hats is red. The fact that girl 2 does not leave the room the second time the ring bells informs girl 1 that her hat is red, for if her hat had been white, girl 2 would have understood that her hat was red, and would have leave the room the second time the bell rang. The same reasoning explains why girls 2 and 3 also leave the room the third time the bell rings.

From these examples, it appears that common knowledge has strong implications for coordination problems. Roughly speaking, a coordination problem is a situation in which each person wants to participate in a joint action only if others participate also. In the “Getting off the bus” example, the two co-workers want to joint their acquaintance and have a drink only if both do it. One way to coordinate could be simply to communicate a message such as “Let’s all participate”. But since each person will participate only if others do, for the message to be successful, it is not sufficient that each person knows about it. Each person must also know that each person knows that each person knows about it... and so on. In other words, people will participate only if it is common knowledge that they all do. In the bus example, the co-workers have to generate common knowledge of the fact that they know about the invitation. If there is uncertainty at some level, like in the case where one of the two co-workers knows that the other knows about the invitation, but does not know whether the other knows that he knows, they do not get off the bus.

Another issue for which common knowledge has strong implications has been illuminated by Aumann [1976]. Aumann showed that rational agents cannot “agree to disagree” about their posterior beliefs, formalized as probability distributions, if they have common prior beliefs. More precisely, if two agents have the same prior probability, and if they have common knowledge of their posterior probability of a given event, then these posterior probabilities must be the same, despite different conditioning information. This result

suggested that asymmetric information had less explanatory power than might be thought: in the absence of differences in prior beliefs, asymmetric information can not explain commonly known differences in posterior beliefs. In particular, Aumann's result has crucial implications for the theoretical analysis of speculation and trade among rational agents. Consider for instance two stock traders who have received contradictory information about the evolution of the price of some stock. Trader A , who believes that the price will go down, offers his stocks to trader B . The deal is concluded, and a handshake makes it common knowledge to both traders. If the fact that traders A and B are willing to exchange is common knowledge to them, then it is also common knowledge to them that A believes the price will go down, and that B believes the price of the stock will go up. Yet this is not possible, according to Aumann's result. To restore the conventional understanding of speculation and trade, one has to assume either that traders are boundedly rational ("noisy traders") or that agents hold different prior probabilities.

Aumann's result gave rise to a vast literature, which addresses basically the same question. To what extent asymmetric information can explain differences in beliefs and decisions?

Motivation and overview of the thesis

This thesis is a contribution to the literature that followed the result of Aumann, which we shall call the *Agreeing to Disagree* literature. It is made of six chapters. The first two chapters introduce background material and survey the literature. The next three chapters are original contributions. The last chapter compares various conditions and assumptions of the three contributions with conditions and assumptions that can be found in the literature.

The notions of knowledge and common knowledge are central to this literature, as well

as their modelling. The first chapter of this thesis is a survey, devoted to a presentation of the model of knowledge used in the Agreeing to Disagree literature and in this thesis. This model, called an Aumann structure, consists of a set of states of the world, and an information partition for each agent. A state of the world is a full description of the world, including objective facts such as “It is raining”, and knowledge facts such as “I know that you know that it is raining”. Information partitions are partitions of the set of states of the world, and are intended to capture the way agents are informed about the world. The framework of this thesis is fundamentally interactive: this thesis deals with the way individuals learn from communicating with each other, and with the inferences that can be made from common knowledge of their actions. The way agents make inferences from observing others’ actions crucially depends on what agents know about the structure of others’ information. Furthermore, when agents make inferences from others’ actions, their knowledge evolve. Yet states of the world are a full description of the world, including agents’ knowledge. Therefore, we make an attempt in Chapter 1 to answer to the two following questions, which are of primary importance for the Agreeing to Disagree literature. *Are individual partitions “common knowledge” to all individuals? Do states of the world evolve when individuals update their information?*

The aim of the second chapter is to review the Agreeing to Disagree literature, and to present how our contributions relate to it. Aumann’s result illuminated a particular property of common knowledge, which is stated as “common knowledge of individual posterior probabilities negates asymmetric information”. In other words, common knowledge of individual posteriors implies that these posteriors do not reflect the differential information that each agent possesses. The first question addressed in the Agreeing to Disagree literature is then the following. *Under what conditions common knowledge of a statistic of individual*

decisions negates asymmetric information? In the Agreeing to Disagree literature, the way agents make their decisions is described by a *decision rule*, which prescribes what decision to make as a function of any information situation agents might be in. Aumann's result provide an answer to this question in the case where the statistic is the identity function, and individual decision rules are posterior probabilities of a given event. McKelvey and Page [1986] show that sufficient conditions are that the statistic is *stochastically regular*, and that individual decision rules are posterior probabilities of some event. Their result implies, for instance, that common knowledge of the mean of individual posteriors implies that individuals have all the same posterior, provided that they have the same prior. In Chapter 3, we investigate the case where decisions may not be posterior probabilities. Suppose for instance that a recruiting committee contemplates hiring somebody for a job. All members have the same preferences, namely they all want to hire to most skilled applicant, but they have received differential information about each applicant's abilities. A poll is realized among members of the commission, who are asked which applicant they think is the most qualified for the job. For each candidate, the percentage of members who think that the candidate should be hired is made public. How will recruiters react to the public announcement of the poll? We show that in that case, they cannot disagree on the candidate they want to hire. More generally, the aim of Chapter 3 is to answer the following question.

1. What conditions should be imposed on the statistic and on individual decision rules to guarantee that common knowledge of the statistic implies that individual decision do not reflect the differential information that each agent possesses, in the case where decisions may not be posterior probabilities?

We show that if the statistic is *exhaustive*, and if individual decision rules are *balanced*

union consistent, which is a weaker requirement than posterior probabilities, then common knowledge of the statistic of individual decisions negates asymmetric information.

It could be argued that common knowledge of individual decisions is a theoretical situation that is not attainable. Therefore, the second question addressed in the Agreeing to Disagree literature deals with how individual decisions can *become* common knowledge in a group of agents. The general framework adopted in this line of research is the following. Agents communicate their decisions according to a protocol upon which they have agreed beforehand, which determines the senders and the receivers of the communication at each date, and update their private information according to what they hear. The question addressed in the Agreeing to Disagree literature is the following. *What conditions should be imposed on the communication protocol and on individual decision rules to guarantee that, eventually, individual decisions become common knowledge to all agents?* Geanakoplos and Polemarchakis [1982] were the first to provide an answer to this question. They showed that if two agents have the same prior probability, and communicate and update their posterior probabilities of some event back and forth, then they will eventually converge to a *consensus*, namely a situation in which they have common knowledge that they have the same posteriors. In Chapter 4, we investigate the case of non-public communication protocols, namely protocols in which agents may privately communicate at some dates. Parikh and Krasucki [1990] were the first to investigate this case. They showed that if individuals follow the same decision rule, and if this decision rule satisfies a *convexity* condition, then communication eventually leads to a consensus. We show in Chapter 4 that Parikh and Krasucki's convexity condition may not apply in some decision spaces, such as finite decision spaces. The typical example is the one of diffusion of innovations. Consider a group of rice producers who have suddenly the opportunity to produce maize. The

suitability of maize for the region's soil is uncertain, and farmers are therefore uncertain about the yield of this new variety of crop. Suppose that each farmer receives private information about the yield of maize, but that they cannot share directly their information with each other, because they must plow their land all day long. The only contact they have with each other is visual: they are able to see their nearest neighbors, but not the neighbors of their neighbors. In particular, they know what kind of crop are cultivated by their neighbors, but not the variety of crop cultivated by the neighbors of their neighbors. Suppose that according to their private information, some farmers decide to adopt maize. What will other farmers do? Is the situation where some farmers produce maize and some others rice a long term equilibrium (in this very simple framework)? We show that from some time on, all farmers must produce the same crop, maize or rice, and that it must be common knowledge to them. More generally, the aim of Chapter 4 is to answer the following question.

2. Under what condition on individual decision rules communication leads to consensus in any fair protocol and in any decision space?

We show that if agents have the same utility function and communicate the action that maximizes their expected utility given their private information, then they will eventually converge to a consensus on their decisions. Contrary to Parikh and Krasucki's convexity condition, our condition applies to any action space.

In this setting, different communication protocols may lead to different outcomes, in terms of consensus decision and of information learned by the agents during the communication process. In particular, it may well be the case that some agent learns more information when communicating according to some protocol α than according to some

protocol β . In the crop example, we know that under certain conditions, all farmers will agree on one unique sort of crop. However, whether this consensus crop will be rice or maize depends on the way farmers are located in the region. Furthermore, a farmer may learn more information if his location is such that he is able to see all other farmers than if he observes only one other farmer. It may well be the case that the most informative protocols are not the same for all agents. Therefore, if agents communicate in order to learn information from each other, they may disagree about the protocol they should use for communicating. For instance, some farmers may fight to buy a certain plot of land, not because this plot is more productive than others, but because it is located on a place where farmers could learn a lot of information from others, about yields of new variety of crops or about ways to produce more efficiently. The aim of Chapter 5 is to answer the following question.

3. If agents communicate so as to learn information, what inferences can they make from common knowledge that some of them disagree about the protocol they should use for communicating?

We show that in that case, the consensus decision must be the same whatever the communication protocol. Our result in the farmers example is as follows. Suppose that it is common knowledge in the region that farmer A wants to exchange his land plot with the one of farmer B , and that farmer B refuses the exchange. Then the consensus crop will be the same whether B exchanges his land plot or not.

All results in the Agreeing to Disagree literature require that individual decision rules satisfy some conditions. In Chapter 2, we define in particular *union consistency*, *convexity* and *weak convexity*, which were used in different settings to guarantee that agents cannot

agree to disagree on their decisions. In Chapter 3 and Chapter 4, we introduce two new conditions called *balanced union consistency* and *argmax*. The aim of Chapter 6 is first to examine how the conditions we introduced relate to convexity, weak convexity and union consistency. Second, we study the consensus properties of each condition in the settings of Chapter 3 and Chapter 4.

Chapter 1

Modelling knowledge

1.1. Introduction

The model of knowledge that is used in this thesis is sometimes called “the standard model of knowledge”, for it is often used in economics and in game theory. Though there are other approaches to modelling knowledge, this one has the advantage of being very tractable in interactive settings, where knowledge about others’ knowledge matters. This model, called an Aumann structure, consists of a set of *states of the world* and of an *information partition* for each agent. Since Savage [1954] and Harsanyi [1967], the state of the world has been the fundamental conceptual tool used for modelling knowledge in game theory and decision theory. A state of the world is a full description of the world, or at least a full description of the relevant facts for the economic problem considered by the modeler. It specifies objective facts, whether past, present or future (*i.e.* the physical environment); it also specifies knowledge facts, that is to say individuals’ knowledge of objective facts, as well as knowledge about individuals’ knowledge, at any level. Information partitions are partitions of the set of states of the world, and are intended to represent the way individuals are informed about the world.

The framework used in this thesis is fundamentally interactive: individuals learn information from communicating with each other, and the analysis focuses on the inferences that can be made by individuals from common knowledge of their mutual decisions. Consider for instance the following situation. An individual, denoted A , has access to two kinds of weather reports: a temperature map, or a precipitation map. If he sees the temperature map, he will know if the next day will be hot or cold. If he sees the precipitation map, he will know if the next day will be rainy or dry. Formally, there are four states of the world: ω_1 =(rainy and hot), ω_2 =(rainy and cold), ω_3 =(dry and hot) and ω_4 =(dry and cold). Imagine that A tells a friend, denoted B , that he intends to go to the beach on next day, and that B knows that A goes to the beach only when he is sure that the weather is hot or dry. What will B learn from the fact that A intends to go to the beach on next day? If B does not know what kind of weather report A saw, B only learns that the weather will not be (rainy and cold). To learn more from A 's decision, B needs to know what kind of information A has. This raises an important question. What do individuals know about the others' information in Aumann structures? In chapters 3, 4, and 5, as well as in the literature in which we cast our analysis, it seems implicit that agents "know" the information partitions of the other agents. What is the status of this assertion? Is it a meta-assumption? Is it, as Aumann wrote, a tautology implicit in the model?

In chapters 4 and 5, we consider a dynamic setting, where individuals communicate with each other and update their information according to what they hear or observe. In this setting, learning explicitly occurs over time. This raises another important question. How to model updating in Aumann structures? In particular, do states of the world evolve over time?

The aim of this chapter is first to present the approach used to model knowledge

in this thesis, namely Aumann structures, and to introduce some notations that will be maintained throughout the thesis. In particular, we present the properties satisfied by individual knowledge in Aumann structures and the modelling of the crucial notion of common knowledge. This modelling raises two questions that are of primary importance for the Agreeing to Disagree literature, and then for this thesis.

1. Do individuals “know” others’ information partitions?
2. Do states of the world change when individuals update their information?

We make an attempt to answer these questions in section 3 of this chapter. To do that, we have go back to the logical foundations of Aumann structures, and to answer the questions of the very nature of states of the world and of information partitions. Though less known by economists, an alternative approach to modelling knowledge is the one traditionally taken in philosophy, epistemic logics and artificial intelligence. Let us call this approach the logic based approach. Going back to at least Hintikka [1962], the logic based approach uses a *logical language*, based on a set of primitive propositions and closed under logical operators. Knowledge is here expressed syntactically, as a modal operator on formulas of the language. The formal model used to determine what formulas are true is a *Kripke structure*. Among other objects, a Kripke structure consists of a possibility relation for each agent, which is intended to capture what states agents consider possible in each state. Aumann structures turned out to be that particular case of Kripke structures in which individual’s possibility relations are equivalence relations. The logic based approach has the advantage to provide an explicit interpretation of states of the world, and to make clear the nature of information partitions.

1.2. Modelling knowledge with Aumann structures

1.2.1. Individual knowledge in Aumann structures

We assume that agents of a group wish to reason about a world that can be described in terms of *objective facts* such as “It is raining”, and of *knowledge facts* such as “I know that it is raining”. A *state of the world* consists of the list of all those facts (objective facts and knowledge facts) that are true in this world. In the spirit of probability theory, the approach to modelling knowledge in economics focuses on *events*, which are subsets of the set of possible worlds Ω . In this thesis, we shall denote¹ $\mathcal{E} \subseteq 2^\Omega$ the set of possible events. More precisely, an event is the “*field*”² of a property, namely the set of states of the world in which this property holds, and events are usually identified with the properties of which those events are the field. For instance, the set of states of the world in which the property “It is raining” is true is identified with the event “It is raining”. Given an event $E \subseteq \Omega$, let us denote p_E some property whose field is E . We say that event E occurs in some state ω if p_E is true at ω . As E is the set of states of the world in which p_E is true, E is said to occur, or to be true at state ω if $\omega \in E$. The empty set corresponds to a contradiction, and the whole set Ω to a tautology, namely a property that is true in every state of the world. Moreover, the inclusion between events corresponds to material implication between properties. As usual, $\neg E$ denotes the negation of event E , namely $\neg E := \Omega \setminus E$.

The standard model used to determine what agents know about the states of the world

¹As it depends on Ω , the set of possible events associated to Ω should be denoted \mathcal{E}_Ω . As we always consider a unique set of states of the world, there is no possible confusion and we omit the subscript.

²We borrow the term of field from logics. In logics, an event is the field of a formula, namely the set of possible worlds in which this formula is true.

is an *information structure*. It consists of a set of states of the world Ω and a *possibility correspondence* P_i for each individual i . A possibility correspondence for an agent is a function $P : \Omega \rightarrow 2^\Omega$, such that $P(\omega)$ is the set of states of the world that the agent conceives as possible at state ω . A possibility correspondence for an agent is intended to capture what worlds the agents considers possible in any given world. At state ω , the agent excludes all states outside $P(\omega)$, and does not exclude any states in $P(\omega)$.

There are two equivalent ways of expressing knowledge in an information structure. The first one is to list the set of events that are guaranteed to occur given the information an individual possesses in some given state. An individual knows an event E at state ω if this event occurs at any state that the agent conceives as possible at ω , namely if $\forall \omega' \in P(\omega), \omega' \in E$. In other words, an agent endowed with an possibility correspondence P is said to know an event E at state ω if and only if:

$$P(\omega) \subseteq E$$

An alternative way of expressing knowledge in an information structure is to represent an individual's knowledge of some event E by enumerating all the possible worlds in which the information that the individual possesses guarantees that E must occur. To do so, one defines individual knowledge operators as functions from the set of events into itself, mapping any event E into the set of states in which the individual knows that E has occurred. Formally, an individual knowledge operator is denoted $K : \mathcal{E} \rightarrow \mathcal{E}$. Given any event E , $K(E)$ is the event "the agent knows E ".

The possibility correspondence of an individual can be constructed on the basis of the knowledge operator, and the knowledge operator can be constructed on the basis of the

possibility correspondence *via* the following relation:

$$K(E) = \{\omega \in \Omega \mid P(\omega) \subseteq E\} \quad (1.1)$$

We denote $\neg K$ the operator “the agent does not know”. It is defined by $\neg K(E) := \Omega \setminus K(E)$ for all event E , and $\neg K(E)$ is the event “the agent does not know E ”. Note that not knowing an event E does not mean knowing that E has not occurred. Whenever an individual knows $\neg E$, namely knows that E has not occurred, this individual does not know E . Formally, $K(\neg E) \subseteq \neg K(E)$:³ if an individual knows that E is false, then he does not know that E is true. However, the converse is not true. Not knowing E does not imply knowing that E has not occurred. Consider the example of an individual endowed with the following possibility correspondence:

$$P(\omega_1) = \{\omega_1\}, P(\omega_2) = \{\omega_2\}, P(\omega_3) = \{\omega_3, \omega_4\}, P(\omega_4) = \{\omega_3, \omega_4\}$$

and consider the event $E = \{\omega_2, \omega_3\}$. The set of states in which the individual knows E is $K(E) = \{\omega_2\}$. Therefore, the set of states in which he does not know E is $\neg K(E) = \{\omega_1, \omega_3, \omega_4\}$. The event $\neg E$ is $\{\omega_1, \omega_4\}$. Therefore, the individual knows $\neg E$ only at state $K(\neg E) = \{\omega_1\}$. The set $K(\neg E)$ is strictly included in the set $\neg K(E)$.

An *Aumann structure* is a particular information structure where each individual possibility correspondence satisfies the three following properties:

For all $\omega \in \Omega$,

$$(PC1) \ \omega \in P(\omega)$$

$$(PC2) \ \forall \omega' \in P(\omega), P(\omega') \subseteq P(\omega)$$

³This is true as long as the state which has actually occurred is always believed possible. $\omega \in K(\neg E) \Rightarrow \omega \in \neg E$. As $K(E) \subseteq E$, one has $\neg E \subseteq \neg K(E)$. Then $K(\neg E) \subseteq \neg K(E)$.

(PC3) $\forall \omega' \in P(\omega), P(\omega) \subseteq P(\omega')$

A possibility correspondence satisfying these three properties defines a partition of Ω , and is called an *information partition*. Typically, we will denote Π the information partition of an agent. Therefore, an Aumann structure consists of a set of states of the world and of an information partition for each agent.

As $K(E)$ is a particular event of the model, the operator $\neg K$ can be applied to $K(E)$. $\neg K(K(E))$ is then the event “the agent does not know that he knows E ”. Is it possible that this event occurs, namely is there some state ω such that $\omega \in \neg K(K(E))$? More generally, what properties of knowledge do we implicitly impose when assuming Aumann structures? The fact that individuals are informed about the set of states of the world by information partitions implies that knowledge operators defined by (1.1) satisfy some properties, which we list in the next section.

1.2.2. *Properties of knowledge in Aumann structures*

In this section, we first list the properties that the individual knowledge operator K_i must satisfy whenever knowledge is partitional. From these axiomatic properties it follows that individuals are aware of all possible contingencies, that they know all the logical implications of their knowledge, that they cannot know things that are false, and that they know what they know, and what they do not know. Second we discuss the implications of these properties. We particularly insist on negative introspection, which has been widely criticized in the literature.

1.2.2.1 *Axiomatic properties of partitional knowledge*

The first two properties of knowledge must be considered apart from the last three.

Indeed, these properties are satisfied in any information structures, not only in partitional ones. They follow from what it means to know something in information structures, and not from the assumption that knowledge is partitional.

A1. $\Omega = K(\Omega)$. This property corresponds to the *axiom of awareness*, or to the *Knowledge Generalization Rule*, for it implies that individuals can identify the entire set of possible worlds: in each state of the world, individuals know Ω . Therefore, agents cannot be surprised by some unforeseen contingency. The reason why this property is satisfied in any information structures is the following. Ω is that particular event that occurs in every state of the world. Therefore, Ω must also occur in any state of the world that the agent conceives as possible in any state of the world, which by definition implies that the agent must know Ω in any state of the world. Therefore, one always has $K(\Omega) = \Omega$. Let us recall that the event Ω corresponds to a tautology, namely a property which is true in any state of the world. The Knowledge Generalization Rule therefore means that individuals always know every property of the model which are necessarily true.

A2. $K(E \cap F) = K(E) \cap K(F)$. This property corresponds to the *distribution axiom*, or *conjunctiveness axiom*, since it allows to distribute knowledge operators over intersection (which represents the conjunction connective). This axiom means that individuals know E and F if and only if they know E and they know F . Knowing the intersection is the intersection of knowing. This axiom implies the following property: if $E \subseteq F$, then $K(E) \subseteq K(F)$, namely whenever an agent knows an event E , which is included in an event F , the agent also knows the event F . The reason why this property is satisfied in any information structures is the following. If the

agent knows E at ω , then E occurs at any state that the agent conceives as possible at ω ($P(\omega) \subseteq E$). If $E \subseteq F$, then F occurs whenever E occurs. Therefore, F must occur at any state that the agent conceives as possible, which means that the agent knows F ($P(\omega) \subseteq E \subseteq F$). Let us recall that inclusion between events corresponds to material implication between properties. Therefore, the distribution axiom implies the property of *logical omniscience*, which states that when an agent knows an event, he also knows all the logical implications of this event.

The following properties are derived from the fact that P satisfies (PC1), (PC2), and (PC3), and are therefore characteristic of Aumann structures.

A3. $K(E) \subseteq E$. This property corresponds to the *truth axiom*. Although agents may not know things that are true, whenever they know an event, this event must be true. If an agent knows E in some state ω , then E occurs in ω . This property of knowledge comes from the fact that the true state of the world is always conceived as possible by individuals when they have information partitions, which corresponds to property (PC1). If an agent knows an event E at state ω , then E must be true in every state that he conceives as possible at ω . As, in particular, ω is a possible state for the at ω , then E must be true at ω .

A4. $K(E) \subseteq K(K(E))$. This property corresponds to *positive introspection* axiom. It means that an agent cannot know an event without knowing he knows it. It comes from the transitivity property of information partitions: $\forall \omega' \in \Pi(\omega), \Pi(\omega') \subseteq \Pi(\omega)$, namely the property (PC2). If the agent considers ω' possible at state ω , then the agent must also consider possible every state ω'' that he would consider possible at state ω' . Suppose that the agent knows an event E at ω , *i.e.* $\Pi(\omega) \subseteq E$. By property

(PC2), we have $\forall \omega' \in \Pi(\omega), \Pi(\omega') \subseteq \Pi(\omega) \subseteq E$, therefore, the agent must know E in any states that he considers possible at ω , which precisely means that the agent knows that he knows E .

A5. $\neg K(E) \subseteq K(\neg K(E))$. This property corresponds to the *negative introspection* axiom. Individuals may not know some events, but in that case they know that they do not know them. If an agent does not know the event E , then he knows that he does not know it. Agents are able to list those events that they do not know at some state. Therefore, this axiom precludes people from ignoring their own ignorance. This follows from the euclidianity property of information partitions: $\forall \omega' \in \Pi(\omega), \Pi(\omega) \subseteq \Pi(\omega')$, namely to the property (PC3). If the agent considers ω' possible at state ω , then he knows at ω all the states that he would conceive as possible at state ω' . Let us show that if an agent does not know that he does not know an event E , then he must know E . Suppose that the agent knows E at some state ω' that he considers possible at ω , *i.e.* $\exists \omega' \in \Pi(\omega), \Pi(\omega') \subseteq E$. By (PC3), $\Pi(\omega) \subseteq \Pi(\omega')$, thus $\Pi(\omega) \subseteq E$, which means that the agent knows E at ω .

Note that A1 is implied by A3 and A5, and that A4 is implied by A3 and A5. These properties together form Bacharach [1985]’s axiomatic system of the partitional model of knowledge. Each of them corresponds to a particular axiom of a well-known axiomatic model in epistemic logics called S5 or KT45. We present Bacharach’s characterization theorem for the sake of consistency with our framework, but this equivalence has been known in epistemic logics at least since Hintikka [1962] or Hugues and Cresswell [1968].

Theorem 1 (Bacharach [1985]) *The individual knowledge operator K satisfies A1 – A5 if and only if there exists an information partition Π such that $K(E) = \{\omega \in \Omega \mid$*

$\Pi(\omega \subseteq E)$ for all $E \subseteq \Omega$.

In the next section, we discuss the properties of partitional knowledge.

1.2.2.2 Discussion of these properties

By imposing the axiom of awareness (A1), it is assumed that in every state of the world, agents are aware of the entire set of possible worlds. This assumption might seem straightforward in simple environments like the one in the *three hats* example, or in simple card games situations. However, numerous realistic examples challenge the plausibility of this axiom. Suppose⁴ that there are two states of nature: either the ozone layer is disintegrating (state H) or it is not (state \bar{H}). Imagine that a decaying ozone layer would emit gamma rays. In state H , scientists would observe gamma rays, then would investigate their cause and deduce that the ozone was disintegrating. They would then be aware of the two possible states H and \bar{H} . In state \bar{H} , scientists would observe no gamma rays, and would not even realize that there *could be* a hole in the ozone layer. Commonly, we realize the current situation, and therefore its possibility as a world, from the moment where this situation changes. For instance, we realize that our computer was making a background noise as soon as we shut down the computer.

By imposing the distribution axiom (A2), one assumes that agents know every implication of everything they know. This property, combined with the fact that agents know all possible contingencies, implies that agents have the property of *logical omniscience*. However, people are obviously not logically omniscient. A common example is that individuals know the rules of chess, but do not know whether they are playing a winning strategy. Lack of logical omniscience may stem from many sources. An obvious one is lack of computational power; for instance, people clearly do not have the computational resources to

⁴This example is from Geanakoplos [1994].

compute a best response strategy in the game of chess. Another cause of lack of logical omniscience is that people may do faulty reasoning. People often happen to believe logical contradictions, for instance when believing ϕ and ψ are true, whereas ψ implies $\neg\phi$.

Therefore, the first two properties of knowledge imply that individuals are very powerful reasoners. However, as we said before, this is not due to the Aumann structure, *i.e.* to the fact that individuals have a partition of Ω . This is due to the definition of knowing an event in information structures, as “an individual knows an event at some state if this event occurs in any state that the agent considers possible at that state.” The last three properties however, are characteristic of Aumann structures, and are usually associated with “rationality regarding knowledge”, which requires that knowledge employed by agents to make their decisions is derived from coherent inferences.

By imposing the truth axiom (A3), one assumes that agents cannot be wrong, namely cannot know things that are not true. If an agent believes that some event has occurred, then this event must have occurred. (A3) has been taken by philosophers to be the major axiom distinguishing knowledge from belief. Although agents may have false beliefs, they cannot know something that is false.

Among the five axioms characterizing Aumann structures, that of Negative Introspection has been, to the best of our knowledge, the one that was most criticized. The following example, first provided by Geanakoplos [1989], is commonly used to illustrate departures from the partitional model of knowledge. Sherlock Holmes and Dr Watson are investigating a crime. From what the local police tell him, Holmes notices that the dog in the garden did not bark that night, and hence concludes that there was no intruder in the garden.

“Is there any other point to which you would wish to draw my attention?”

“To the incident of the dog in the night-time.”

“The dog did nothing in the night-time.”

“That was the curious incident”, remarked Sherlock Holmes.

Doyle [1901]

Geanakoplos’ interpretation of the “curious incident” example is the following. When Watson says that “The dog did nothing in the night-time”, he seems not to be aware that “not barking” is actually *doing something*. Formally, reducing the set of possible worlds to two states, B standing for the state in which the dog did bark and \bar{B} for the state in which the dog did not bark, Geanakoplos defines possibility correspondences of Holmes and Watson as follows:

$$\left\{ \begin{array}{l} P_H(B) = \{B\} \\ P_H(\bar{B}) = \{\bar{B}\} \end{array} \right. \text{ and } \left\{ \begin{array}{l} P_W(B) = \{B\} \\ P_W(\bar{B}) = \{B, \bar{B}\} \end{array} \right.$$

Clearly, Holmes’ possibility correspondence satisfies the three properties ensuring that it defines a partition of $\{B, \bar{B}\}$. Watson’s possibility correspondence however does not satisfy the property (PC3) corresponding to the negative introspection axiom. Indeed, $B \in P_W(\bar{B})$, but $P_W(\bar{B}) \not\subseteq P_W(B)$. In the state where the dog did not bark, Watson does not know that “the dog did bark that night”, and does not know he does not know either.

There are different ways of interpreting this story. We can imagine that when Watson says that the dog did nothing, he understands that the dog did not bark. However, he somehow does not come up with the inference that there was no intruder. Watson and Holmes both received the same information, and yet, contrary to Holmes, Watson did not perceive the fact “The dog did not bark” as a clue. Therefore, this story can also be modelled as an example of violation of the logical omniscience property. Watson knows that the dog did not bark, but he does not conclude anything from it, in particular he does not conclude that there was no intruder in the garden. However, recall that the logical

omniscience property is *de facto* satisfied in information structures. Therefore, one should use another model of knowledge.

Li [2006, p 2] proposes an alternative explanation of this story. “*Watson is unaware of the possibility that there was no intruder, and hence fails to recognize the factual information “there was no intruder” contained in the message “the dog did not bark”. Had someone asked Watson, “Could there have been an intruder in the stable that night?”, he would have recognized his negligence and replied “Of course not, the dog did not bark!”.*” This explanation suggests that one could also explain Watson’s failure by a violation of the axiom of awareness. Here again, one should use another model of knowledge than information structures, in which the axiom of awareness is always satisfied.

We do not claim that negative introspection is not a strong assumption. However, we think that the stronger assumptions that are made on individual knowledge are the axiom of awareness (A1) and the one of logical omniscience (A2). Yet these two axioms are precisely the ones that are not due to the assumption that knowledge is partitional.

We may wonder what would be the consequences of relaxing the Negative Introspection assumption. An interesting implication deals with the value of information. We know since Blackwell [1953] that in non-interactive settings, having more information is always an advantage for a decision maker who maximizes an expected utility. Consider a decision maker who has to choose an action whose payoff depends on the state of the world. Formally, denote \mathcal{A} the set of available actions and $U : \mathcal{A} \times \Omega \rightarrow \mathbb{R}$ the decision maker’s payoff. The decision maker has a prior probability P over Ω , and takes the action that maximizes his expected utility *ex interim*: the state of the world ω occurs, the decision maker is informed of $\Pi(\omega)$ and then takes the decision that maximizes his expected utility conditionally to $\Pi(\omega)$. Denote $f_{\Pi}(\omega)$ the action chosen by the decision maker at state ω ,

given his partition Π , formally $f_{\Pi}(\omega) := \operatorname{argmax}_{a \in \mathcal{A}} E[U(a, \cdot) \mid \Pi(\omega)]$.

Proposition 1 *Let Π, Π' be two information partitions. If Π is finer than Π' , then the ex ante expected payoff yielded by Π is larger than that yielded by Π' . Formally, $E[U(f_{\Pi}(\cdot), \cdot)] \geq E[U(f_{\Pi'}(\cdot), \cdot)]$.*

With non-partitional information however, having more information may not be better for the decision maker. Consider the following example, drawn from Rubinstein [1998]. There are three equally likely states of the world, denoted $\omega_1, \omega_2, \omega_3$. A seller offers a risk-neutral decision maker the following bet: the decision maker gets 3 euro if state ω_2 occurs, and -2 euro if states ω_1 or ω_3 occurs. If the decision maker has no additional information about the states of the world, the best option for him is to reject the bet, as his expected payoff of accepting the offer is $-1/3$, whereas it is 0 if he rejects the offer. To persuade the decision maker to accept the offer, the seller gives the decision maker the following “bonus” information. If the state is ω_1 , the decision maker is told that ω_3 has not occurred. If the state is ω_2 , he is told that neither ω_1 nor ω_3 have occurred. Finally, if the state is ω_3 , the decision maker is told that ω_1 has not occurred. Assuming that he does not understand the rule by which the bonus information is given, the decision maker is now endowed with the following possibility correspondence.

$$P(\omega_1) = \{\omega_1, \omega_2\}, P(\omega_2) = \{\omega_2\}, P(\omega_3) = \{\omega_2, \omega_3\}$$

Such a possibility correspondence violates only property (PC3), which corresponds to the negative introspection axiom. According to it, the best option of the decision maker is to accept the offer at any state. His ex ante expected payoff is then $-1/3$. Therefore, being more informed is not necessarily an advantage when knowledge is not partitional.

1.2.3. Common knowledge in Aumann structures

Perhaps first introduced by Lewis [1969], the informal definition of *common knowledge* is stated as follows: an event is said to be common knowledge in a group of agents if all individuals in the group know it, all individuals know that all individuals know it, and so on *ad infinitum*. Aumann [1976] was the first to provide a formal characterization of the notion of common knowledge.

Like individual knowledge, common knowledge is a property which is true in some states and false in others. With individual knowledge operators, Lewis' definition of common knowledge is as follows: an event E is common knowledge at state ω in a group of agents \mathcal{N} if for all integer n , and all sequence (i_1, \dots, i_n) such that $i_k \in \mathcal{N}$, one has $\omega \in K_{i_1}(K_{i_2}(\dots K_{i_n}(E)\dots))$. In Aumann's words, E is common knowledge at ω if it contains all states ω' that are *reachable* from ω , given the individuals' information partitions. Therefore, even if there are finitely many agents and if individual information partitions are finite, one has to check infinitely many conditions to know whether E is common knowledge at ω . Aumann [1976] showed the equivalence of this definition with another one, which is more tractable. Let us first define the meet of a set of partitions.

Definition 1 (Meet of partitions) *The meet of partitions Π_1, \dots, Π_n , denoted M , is the finest common coarsening of these partitions, namely the finest partition such that $\forall i, \forall \omega, \Pi_i(\omega) \subseteq M(\omega)$.*

In other words, the meet of a collection of partitions $(\Pi_i)_i$ is the finest partition whose cells are a union of cells of each Π_i . Aumann [1976] showed that an event E is common knowledge at state ω if E includes that member of the meet of the individuals' partitions that contains ω .

Proposition 2 (Aumann [1976]) *Let \mathcal{N} be a group of agents, each agent $i \in \mathcal{N}$ being endowed with an information partition Π_i , and let M denote the meet of partitions $(\Pi_i)_{i \in \mathcal{N}}$. An event E is common knowledge at state ω iff $M(\omega) \subseteq E$.*

The meet of individual partitions is therefore the partition of common knowledge in the group of agents. According to this definition, finitely many steps are sufficient to check whether an event is common knowledge, in the case where the set of agents and information partitions are finite.

Milgrom [1981] provided a list of characteristic properties of common knowledge. He defined the common knowledge operator as a function $C : \mathcal{E} \rightarrow \mathcal{E}$ such that for all $E \in \mathcal{E}$, $C(E) = \{\omega \in \Omega \mid E \text{ is common knowledge at } \omega\}$. Consider two events E, E' and the following four properties:

$$(P1) \quad C(E) \subseteq E$$

$$(P2) \quad \forall \omega \in C(E), \forall i \in \mathcal{N}, \Pi_i(\omega) \subseteq C(E)$$

$$(P3) \quad E \subseteq E' \Rightarrow C(E) \subseteq C(E')$$

$$(P4) \quad [\forall i, \forall \omega \in E, \Pi_i(\omega) \subseteq E] \Rightarrow E = C(E)$$

Condition (P1) is analogous to the Truth axiom. It asserts that an event E can be common knowledge only if E actually occurs. Condition (P2) holds that if E is common knowledge, then every agent knows that E is common knowledge. Condition (P3) is analogous to logical omniscience: whenever E is common knowledge, then any logical consequence of E is also common knowledge. Condition (P4) deals with particular events called *public events*. An event is said to be public in a group of agents if every agent

knows this event whenever it occurs. Condition $(P4)$ states that public events are common knowledge whenever they occur.

Using Aumann's definition of common knowledge, Milgrom showed that common knowledge is characterized by these four properties.

Theorem 2 (Milgrom [1981]) *There is a unique function C satisfying $(P1) - (P4)$ and it is given by $C(E) = \{\omega \in \Omega \mid M(\omega) \subseteq E\}$.*

As common knowledge can be represented by a particular partition of the set of states of the world, its properties must be the same as that of individual knowledge: awareness, distribution, truth, positive and negative introspection. $(P1)$ and $(P3)$ correspond to the truth and logical omniscience properties. $(P2)$ and $(P4)$ directly imply the positive introspection axiom, and $(P2)$ together with the fact that individual knowledge is partitional imply the negative introspection axiom. As for individual knowledge, negative introspection and logical omniscience imply awareness.

1.3. Controversial issues

Aumann structures are often used to model knowledge in economics and game theory. However, the implicit assumptions of this approach to modelling knowledge have been widely criticized in the literature. One of these criticisms deals with the cognitive abilities that are assumed to have individuals in Aumann structures, which we discussed in the former section. In this section, we aim to present two controversial issues about Aumann structures, which are of primary importance for the Agreeing to Disagree literature.

The Agreeing to Disagree literature deals with the implications of common knowledge of individual decisions. In chapter 3, we consider the particular issue of common knowledge

of an aggregate of individual decisions. How can we interpret the fact that individual decisions, or more generally, events, are common knowledge? Consider two agents A and B endowed with the following partitions:

$$\Pi_A = \{\omega_1\}\{\omega_2, \omega_3\}\{\omega_4, \omega_5\}$$

$$\Pi_B = \{\omega_1, \omega_2\}\{\omega_3\}\{\omega_4\}\{\omega_5\}$$

and imagine that A decides to buy a car in states ω_1, ω_2 and ω_3 , and not to buy it otherwise.

Denoting F the event “ A intends to buy the car”, we have $F = \{\omega_1, \omega_2, \omega_3\}$.

The meet of these two partitions is $M : \{\omega_1, \omega_2, \omega_3\}\{\omega_4\}\{\omega_5\}$. By Aumann’s characterization, any event containing $\{\omega_1, \omega_2, \omega_3\}$ is common knowledge in ω_1 , for instance. Therefore, it is common knowledge at state ω_1 that A intends to buy the car. How can we interpret this assertion? At state ω_1 , B knows that A intends to buy the car as it is the case in any state that he conceives as possible at state ω_1 , namely $\Pi_B(\omega_1) = \{\omega_1, \omega_2\} \subseteq F$. The event “ B knows F ”, namely the set of states in which B knows that A intends to buy the car is $\{\omega_1, \omega_2, \omega_3\}$. As a consequence, the event “ B knows F ” is identical to the event F . A knows that he intends to buy a car at state ω_1 , as $\Pi_A(\omega_1) = \{\omega_1\} \subseteq F$. Can we conclude that A also knows that B knows that A intends to buy the car? How comes that A knows that the event F is also the event B knows F ? It seems that we must suppose that A understands B ’s partition. As F is common knowledge at state ω_1 , the reasoning “ A knows that B knows that A knows etc”, must be iterated *ad infinitum*. Therefore, it seems that underlying Aumann’s definition of common knowledge is the assumption that A and B understand each other’s partition and that this understanding is somehow “common knowledge”. This raises two related questions of primary importance for the Agreeing to Disagree literature.

1. Are individual partitions “common knowledge” to all individuals?

2. Is “common knowledge” of individual partitions necessary for the meet to be the partition of common knowledge?

In chapters 4 and 5, we deal with a dynamic setting where agents learn from each other by communicating. We consider for instance the following situation. Suppose that B intends to buy a car at states ω_1 and ω_2 , but not at ω_3 , and that he tells A about his intention. A will learn something, precisely A will learn to distinguish between states ω_2 and ω_3 . How to model learning processes in Aumann structures? As A learned something from B 's message, A 's knowledge has evolved. Yet states of the world must be a complete description of the world, including A 's knowledge. This raises another important question about Aumann structures.

3. Do states of the world evolve when individuals update their information?

To answer these questions, one has to understand the very nature of states of the world, and their relation with information partitions. A first answer in the economic literature was given by Aumann in 1976: any aspect of the environment about which individuals might be uncertain should be captured in the description of the states in the model: “Included in the full description of a state [...] of the world is the manner in which the information is imparted to the two persons.” (p 1237). This has given rise to a philosophical questioning about the relation between states and partitions in Aumann structures. For instance, Fagin *et al.* [1999] wrote: “If we think of a state as a complete description of the world, then it must capture all the agents’ knowledge. Since the agents’ knowledge is defined in terms of the partitions, the state must also include a description of the partitions. This seems to lead to circularity, since the partitions are defined over the states, but the states contain a description of the partitions.” (p 332). Partly in response to these circularity concerns,

Fagin *et al.* [1991, 1992, 1999] provided an alternative approach to modelling knowledge, in the line of Harsanyi [1968] and Mertens and Zamir [1985]. This approach intends to capture the idea of a complete description of an agent’s uncertainty by constructing an infinite information partition hierarchy, where each level in the hierarchy simply describes which members of the previous hierarchy the agent considers possible. We do not discuss the hierarchical approach here. Rather, we think that to understand what it means for a state to be a “full description of the world”, one has to use the formal language of epistemic logics.⁵

1.3.1. Logical foundations of Aumann structures

The idea of a formal logical analysis of reasoning about knowledge goes back at least to Hintikka [1962]. To carry out complicated reasoning about knowledge, modal logic uses a particular language called a *syntax*. To describe the language, one starts with a nonempty set P of primitive propositions. These primitive propositions stand for objective facts about the world such as “It is raining” or “I take my umbrella”. To express more complicated statements like “It is raining and I am not taking my umbrella”, one uses logical connectives of negation: \neg and conjunction: \wedge . To express statements like “I know that it is raining”, one introduces *modal operators* k_1, \dots, k_n , such that $k_i p$ reads “Agent i knows the proposition p ”. We also use the standard abbreviations $\phi \vee \psi$ (ϕ “or” ψ) for $\neg(\neg\phi \wedge \neg\psi)$, $\psi \rightarrow \phi$ (ψ “implies” ϕ) for $\neg\psi \vee \phi$, and $\phi \leftrightarrow \psi$ (ϕ “is equivalent” to ψ) for $(\psi \rightarrow \phi) \wedge (\phi \rightarrow \psi)$.

We denote Φ the set of *well-formed formulae*, which is the closure of the set of primitive propositions under negation, conjunction, and the modal operators $(k_i)_i$. For instance,

⁵Most of the material about epistemic logic was found in Fagin *et al.* [1995].

considering the primitive propositions r_A , r_B , w_A , respectively standing for *Alice's hat is red*, *Bob's hat is red*, and *Alice's hat is white*, we can express the complicated sentence “Alice does not know whether her hat is red and knows that Bob knows that his hat is red” quite simply in modal logic, with the formula:

$$\neg k_A r_A \wedge k_A k_B r_B$$

A *syntax* consists of a set of well-formed formulas, of inference rules and of an axiomatic system. Once the syntax has been defined, one needs *semantics*, which is a formal model that can be used to determine whether a given formula is true or false. One possible approach is to use a formalization in terms of *Kripke structures*. A Kripke structure for n agents is a tuple $M = (\Omega, \pi, \mathcal{K}_1, \dots, \mathcal{K}_n)$, where Ω is the set of *states of the world*; π is an *interpretation function*, that associates with each state in Ω a truth assignment to the formulae in P , namely $\pi(\omega) : P \rightarrow \{\text{true}, \text{false}\}$ for each $\omega \in \Omega$; finally, for each agent i , \mathcal{K}_i is a binary relation on Ω , that is a set of pairs of elements of Ω . The truth assignment $\pi(\omega)$ tells us whether each primitive proposition is true or false in state ω , and in structure M . The binary relations \mathcal{K}_i capture the possibility relations between states according to agent i : $(\omega, \omega') \in \mathcal{K}_i$ if agent i considers state ω' possible, given his information in state ω . Under the axioms of the system $S5$ or $KT45$, (which is “analogous” to the axiomatic system presented in section 1.2.2 for knowledge operators k_i), the possibility relation \mathcal{K}_i is an equivalence relation, and can be equivalently represented by a partition Π_i , with $\Pi_i(\omega) = \{\omega' \in \Omega \mid (\omega, \omega') \in \mathcal{K}_i\}$ for all ω .

We now define what it means for a formula ϕ to be true in world ω in a structure M , which is denoted $(M, \omega) \models \phi$ and can be read as “ ϕ is true at (M, ω) ”. The relation \models is defined by induction on the structure of each formula, because the notion of a formula is defined inductively.

We start with the primitive propositions:

$$(1) (M, \omega) \models p \text{ iff } \pi(\omega)(p) = \text{true} \text{ and } (M, \omega) \not\models p \text{ iff } \pi(\omega)(p) = \text{false}$$

A conjunction $\phi \wedge \phi'$ is true if both ϕ and ϕ' are true:

$$(2) (M, \omega) \models \phi \wedge \phi' \text{ iff } (M, \omega) \models \phi \text{ and } (M, \omega) \models \phi'$$

A negated formula $\neg\phi$ is true if ϕ is not true:

$$(3) (M, \omega) \models \neg\phi \text{ iff } (M, \omega) \not\models \phi$$

The last condition is the one which defines what it means for an agent to know a formula. Agent i knows ϕ at state ω of structure M if ϕ is true in every state that i considers possible in state ω :

$$(4) (M, \omega) \models k_i\phi \text{ iff } (M, \omega) \models \phi \forall \omega' \text{ such that } (\omega, \omega') \in \mathcal{K}_i$$

The term “full description of the world” is now well understood. We can think about a state as the set of formulas that are true at that state. What a state can be, namely what set of formulas can be a full description of the world, is determined by the axioms of the model. For instance, formulas ϕ and $\neg\phi$ cannot be part of a description together. Additional constraints on the information structure imply, for instance, that $k_i\phi$ and $\neg k_i[k_i\phi]$ cannot belong to the same description.

Let us now go back to the circularity issue. Fagin *et al.* wrote the circularity comes from the fact that “partitions are defined over the states, but the states contain a description of the partitions”. By examining the logical foundations of Aumann structures, it is clear that states do not contain a description of individual partitions. It is true that states “must capture all the agents’ knowledge”, but the agents’ knowledge is not defined in terms of

partitions. Knowledge focuses on formulas, and states of the world describe what formulas are known by agents in that world. Information partitions are representations of what states agents conceive as possible in each state.

We now make an attempt to answer the first two questions about what individuals know about others' knowledge

1.3.2. *What do individuals know about the others' knowledge?*

The first question is whether individual information partitions are “common knowledge” to all individuals. Let us first examine whether individuals “know” the others' information partitions.

Possibility relations, and therefore information partitions, are completely determined by individual knowledge operators with relation (4). Therefore, wondering whether individuals know others' partitions in Aumann structures amounts to wondering whether they know others' knowledge operators in the syntactic formalism. Yet the statement “knowing a knowledge operator” makes no sense in syntax, it is not a well-formed formula. Indeed, knowledge operators operate on formulas, taking each formula ϕ to another formula $k_i\phi$. One can express that j knows that i knows ϕ with the statement $k_j[k_i\phi]$, as $k_i\phi$ is a particular formula. However, one cannot express that j knows k_i , as k_i itself is not a formula. Let us go back to the example of page 32. We addressed the question of how come that A knows that the event F is also the event “ B knows F ”. Looking at the logical foundations of a state of the world, we realize that it is like wondering how A knows that F is the event “ A intends to buy a car”. States describe primitive propositions (“ A intends to buy a car”), as well as formulas using knowledge operators (“ B knows that A intends to buy a car”). Knowing about others' knowledge of formulas is treated as knowing about

primitive propositions: it is part of the description of the states.

Therefore, it makes no sense, to the best of our understanding, wondering whether individuals know other's information partitions, since partitions are just representation tools, used by the modeler. Interpretation functions and information partitions play the same role: the first allow to represent which primitive propositions are true in each state, and the second allow to represent which formulas are known in each state.

The second question is whether “common knowledge” of individual partitions is necessary for the meet to be the partition of common knowledge. Like individual knowledge, common knowledge of formulas is part of the description of the states. Let us introduce the common knowledge operator C . The sentence $(M, \omega) \models C\phi$ states that ϕ is common knowledge in state ω and in structure M . In a number of papers, Fagin, Halpern, Moses and Vardi, as well as Lismont and Mongin investigated the axiomatization of common knowledge by providing suitable axioms on the common knowledge operator C . A review of these theorems can be found in Lismont and Mongin [1994]. It has been shown that in an Aumann structure $M = (\Omega, (\Pi_i)_i)$, $(M, \omega) \models C\phi$ if and only if the set $\{\omega' \in \Omega \mid (M, \omega') \models \phi\}$ includes that member of the meet of partitions Π_i 's that contains ω . Therefore, the meet of individual partitions represents common knowledge in an Aumann structure, as well as individual partitions represent individual knowledge in an Aumann structure. It is a representation and does not require any assumptions such as “common knowledge” of individual partitions.

1.3.3. *Do states of the world change over time?*

We now turn to the third question, which is whether states of the world evolve when agents update their information. In chapters 4 and 5, we consider a dynamic setting where

agents communicate with each other and update their information according to what they hear. In line with the Agreeing to Disagree literature, we model information updating as follows. Let us define a message as a function $m : \Omega \times \mathbb{N} \rightarrow \mathcal{M}$, where $m(\omega, t)$ is the message heard by an agent at date t and state ω . (We do not need to go further into details for the issue we consider in this chapter.) When hearing the message $m(\omega, t)$, agents update their information by eliminating the states in which they would have heard another message. Denote Π^t the information partition of an agent at date t . Formally, the agent updates his partition according to the following revision rule:

$$(\mathbf{RR}) \quad \Pi^{t+1}(\omega) = \Pi^t(\omega) \cap \{\omega' \in \Omega \mid m(\omega', t) = m(\omega, t)\}$$

The agent does not have the same knowledge at date $t + 1$ as at date t . We saw in the previous section that states are a full description of the world, including a description of agents' knowledge. As a consequence, can we consider that states of the world are still the same? In this thesis, we use implicit states of the world, that is to say we do not explicitly describe the list of formulas that are true in each state. However, we may wonder whether the revision rule **RR** that we use in the Agreeing to Disagree literature makes sense, namely whether updated partitions effectively represent *new* knowledge of the *same* uncertainty. We shall show that, to the best of our understanding, states of the world do not evolve if they explicitly describe knowledge at each date, *i.e.* actual knowledge and updated knowledge.

Consider a simple example. Let p denote the primitive proposition “It is raining”, and $\neg p$ the proposition “It is not raining”. Consider two agents, endowed with the following partitions of $\Omega = \{\omega_1, \omega_2\}$:

$$\Pi_A = \{\omega_1\}\{\omega_2\}$$

$$\Pi_B = \{\omega_1, \omega_2\}$$

and assume that the proposition p is true at state ω_1 , and is false at state ω_2 . Let us list some formulas that are valid in states ω_1 and ω_2 :

$$(M, \omega_1) \models p ; k_{Ap} ; \neg k_{Bp} \wedge \neg k_B \neg p$$

$$(M, \omega_2) \models \neg p ; k_{A\neg p} ; \neg k_{Bp} \wedge \neg k_B \neg p$$

In words, at ω_1 , p is true, A knows p and B does not know p and does not know $\neg p$. At state ω_2 , p is false, A knows that p is false and B does not know whether p is true or false. A is able to distinguish between the state in which it is raining and the state in which it is not. Imagine that A says to B whether it is raining or not. By the revision rule (RR), B 's partition changes, and the model becomes:

$$\Pi_A = \{\omega_1\}\{\omega_2\}$$

$$\Pi_B = \{\omega_1\}\{\omega_2\}$$

Let us list some formulas that are now valid in states ω_1 and ω_2 :

$$(M, \omega_1) \models p ; k_{Ap} ; k_{Bp}$$

$$(M, \omega_2) \models \neg p ; k_{A\neg p} ; k_{B\neg p}$$

Now, in world ω_1 , p is true and A and B both know it, and in world ω_2 , p is false and A and B both know it. As the list of formulas that are valid in each state are not the same before and after A 's message, ω_1 and ω_2 do not represent the same states of the world by definition.

In this example, the problem comes from the fact that states only describe static knowledge. To treat the case of a dynamical setting in which agents revise their knowledge, one has to consider states that describe initial knowledge and revised knowledge. Bonanno [2004] and Board [2004] propose a unifying framework for static belief and belief revision. The idea is to augment the language with two operators for each agent, I_i and k'_i . The

interpretation of each operator is as follows: $k_i\phi$ means that i initially knows ϕ (at time 0), $I_i\phi$ means that i is informed that ϕ (between time 0 and time 1), $k'_i\phi$ means that i knows ϕ at time 1, namely after revising his knowledge in light of the information received. The description of each state in the previous example should then be as follows:

$$(M, \omega_1) \models p ; k_1p ; \neg k_2p \wedge \neg k_2\neg p ; I_2p ; k'_1p ; k'_2p$$

$$(M, \omega_2) \models \neg p ; k_1\neg p ; \neg k_2p \wedge \neg k_2\neg p ; I_2\neg p ; k'_1\neg p ; k'_2\neg p$$

Now, the description of each ω is the same before and after A 's message.

To conclude, states of the world do not evolve when agents update their information if one consider sufficiently rich states of the world, which describe individuals' knowledge at each possible date.

1.4. Conclusion

The aim of this chapter was first to present Aumann structures, which are used to model knowledge in the Agreeing to Disagree literature. In this thesis, we model situations in which agents update their private information from communicating with each other. The way agents make inferences from observing others' actions depends on two things. First, it depends on whether they know how others' actions relate to their private information. Second, it depends on whether they understand the structure of others' information. In the example given in the introduction, B knows that his friend goes to the beach only when he knows that the weather is hot or dry. However, he does not know what kind of weather report his friend saw. In the Agreeing to Disagree literature, it seems implicit that agents "know" others' information partitions. The second aim of this chapter was therefore an attempt to make clear whether "common knowledge" of partitions is a meta-assumption

or a tautology implicit in the model. Going back to the logical foundations of Aumann structures, we saw that what agents know is part of the description of the states, and that information partitions are only a useful representation of individual possibility relations. Therefore, it makes no sense, to the best of our knowledge, considering the possibility that agents could have knowledge of a representation tool. Another important question that often arises in the agreeing to disagree literature is whether states of the world evolve when agents update their information. We saw that in dynamic settings, states of the world must describe knowledge at each date, as well as the arrival of new information.

Chapter 2

A review of the *Agreeing to Disagree* literature

2.1. Introduction

A formal and tractable definition of common knowledge was introduced into the economics literature by Aumann [1976]. Assuming that individual knowledge is described by a partition of the set of states of the world, he showed that common knowledge in a group of agents can be represented by a particular partition, derived from each individual partition. In this framework, Aumann showed a mathematically almost trivial but yet very powerful result: if rational agents have the same prior probability, and if their posterior probabilities of some given event are common knowledge at some state, then these posteriors must be the same in that state, despite different conditioning information. In other words, rational agents cannot *agree to disagree*. This result suggested that asymmetric information had less explanatory power than might be thought: in the absence of differences in prior beliefs, asymmetric information could not explain commonly known differences in posterior beliefs. In particular, Aumann's result has crucial implications in the theoretical analysis

of speculation and trade among rational agents. Consider for instance two stock traders who have received contradictory information about the evolution of the price of some stock. Trader A , who believes that the price will go down, offers his stocks to trader B . The deal is concluded, and a handshake makes it common knowledge to both traders. If the fact that traders A and B are willing to exchange is common knowledge to them, then it is also common knowledge to them that A believes the price will go down, and that B believes the price of the stock will go up. Yet this is not possible, according to Aumann's result. To restore the conventional understanding of speculation and trade, one has to assume either that traders are boundedly rational ("noisy traders") or that agents hold different prior probabilities.

Aumann's result gave rise to a literature that we shall call the *Agreeing to Disagree* literature. This thesis presents three results that are part of it. The aim of this chapter is to review the Agreeing to Disagree literature, before introducing our contributions to it. Other treatments include Bonanno and Nehring [1997], Geanakoplos [1994] and Samuelson [2004]. Bonanno and Nehring provided a detailed survey of the Agreeing to Disagree literature, as well as a deep analysis of the justifications and the consequences of the common prior assumption. Geanakoplos wrote a rather technical survey of the implications of common knowledge for economic behavior, whereas Samuelson examines knowledge modelling and its role in economics in a less technical survey. Geanakoplos and Samuelson's reviews are broader in scope and are therefore less exhaustive on the Agreeing to Disagree literature strictly speaking than that of Bonanno and Nehring or than ours.

In his agreement result, Aumann makes the assumption that individuals have the same prior probability, and shows under this assumption that asymmetric information cannot explain commonly known differences in posterior probabilities. Even without the common

prior assumption, common knowledge of individual posteriors implies that these posteriors would have been the same, had agents conditioned on the basis of the public information. In other words, common knowledge of individual posteriors implies that posteriors do not reflect the differential information that each agent possesses. This basic property of common knowledge, identified thanks to Aumann's result, is that *common knowledge of individual posteriors negates asymmetric information*. Aumann's result gave rise to a line of research that studies the conditions under which common knowledge of individual decisions negates the explanatory power of asymmetric information. It could be argued that common knowledge of individual decisions, and common knowledge more generally, is a theoretical situation that is not attainable, which would diminish the importance of Aumann's result. Another line of research has then studied the conditions of emergence of common knowledge of individual decision rules.

The Agreeing to Disagree literature addresses the two following questions:

1. Under what conditions common knowledge of a statistic of individual decisions negates asymmetric information?
2. Under what conditions individual decisions might become common knowledge in a group of agents?

In section 3, we review some of the results that provided an answer to the first question. The way agents make their decisions is described by a *decision rule*, which prescribes what action to make as a function of any information situation they might be in. These results give conditions on individual decision rules and on the statistic of individual decisions which are sufficient to guarantee that common knowledge of the statistic implies that all decisions are made on the basis of the same information. In Aumann's result, these conditions are

that the statistic is the identity function, and that individual decision rules are posterior probabilities of some event. If, moreover, individuals follow the same decision rule (as in Aumann's results where agents have the same prior), then all individuals must take the same decision, a situation which is referred to as *consensus*. Most authors presented their results as being about how common knowledge might imply consensus, assuming commonness of decision rules. We emphasize that the very contribution of these results is rather to provide conditions under which common knowledge negates asymmetric information, which does not require that agents follow the same decision rules. However, commonness of decision rules plays a crucial role in the answer to the second question addressed in the literature.

In section 4, we review some results that provided conditions under which communication might create common knowledge of individual decisions. The setting of these results is the following. Agents communicate their decisions according to a protocol upon which they have agreed beforehand. The communication protocol determines the senders and the receivers of the communication at each date. We distinguish between *public* and *non-public* protocols, to illuminate the role of commonness of decision rules. In public protocols, all agents are receivers of the communication at each date. Provided a *fairness* condition on the protocol, communication according to a public protocol leads to common knowledge of individual decisions without any restriction on decision rules. In non-public protocols however, individual decisions may fail to ever become common knowledge if agents follow different decision rules. We show that in non-public protocols, common knowledge of individual decisions emerge *via* the consensus, and then requires commonness of decision rules.

Among the criticisms addressed to the Agreeing to Disagree literature, two criticisms

could particularly be addressed to this thesis. We discuss these criticisms in section 5.

2.2. Agreeing to disagree

In this section, we present the seminal result of Aumann [1976]. Take the example of the two traders in a formal setting. Suppose that the set of states of the world allows to consider the event E “*The price of the stock will go up*”, and the event \bar{E} “*The price of the stock will go down*”. Traders A and B are endowed with information partitions Π_A and Π_B , and share a common prior probability P over Ω . Suppose that the two traders’ behavior is the following. They buy the stock if they believe that its price will go up with a probability larger than $1/2$. They sell the stock if they believe that its price will go down with a probability smaller than $1/2$. If they believe that the price will go up or down with the same probability, they decide not to trade.

When the state of the world ω occurs, each trader privately receives information $\Pi_A(\omega)$ and $\Pi_B(\omega)$. Both update their probability that the price will go up, which become $P(E | \Pi_A(\omega))$ and $P(E | \Pi_B(\omega))$. Assume that A accepts to sell the stock, and that B accepts to buy it. According to the traders’ behavior rule, it must be the case that $P(E | \Pi_A(\omega)) < 1/2$ and $P(E | \Pi_B(\omega)) > 1/2$. We assume that when the deal is concluded, the fact that A and B accept the deal is made common knowledge to both of them, for instance by a handshake. Therefore, the event “*A thinks the price will go up with a probability strictly smaller than $1/2$ and B thinks the price will go up with a probability strictly larger than $1/2$* ” is common knowledge to A and B at state ω . Denoting M the partition of common knowledge among A and B , we have:

$$M(\omega) \subseteq \{\omega' \in \Omega \mid P(E | \Pi_A(\omega')) < 1/2 \text{ and } P(E | \Pi_B(\omega')) > 1/2\}$$

By Aumann's definition, the partition M is such that for all ω , $\Pi_A(\omega) \subseteq M(\omega)$ and $\Pi_B(\omega) \subseteq M(\omega)$. Therefore, each cell of M is a union of cells of Π_A , and a union of cells of Π_B . As these unions are necessarily disjoint, $P(E | M(\omega))$ is a convex combination of those values $P(E | \Pi_A(\omega'))$ such that $\Pi_A(\omega') \subseteq M(\omega)$. Yet for all $\omega' \in M(\omega)$, we have $P(E | \Pi_A(\omega')) < 1/2$. Therefore, $P(E | M(\omega)) < 1/2$. As $M(\omega)$ is also a disjoint union of cells of Π_B , the same reasoning implies that $P(E | M(\omega)) > 1/2$.

Therefore, given some event E and some value a , it cannot be common knowledge among two agents that one has a posterior probability of E strictly larger than a , and the other a posterior probability of E strictly smaller than a . This result is a slight generalization of Aumann's result, which is known as "*Rational agents cannot agree to disagree*". Notice that the word "agree" plays two different roles in the phrase "agree to disagree": "agree" refers to common knowledge, while "disagree" refers to reaching different decisions.

Theorem 1 (Aumann [1976]) *Consider two agents A and B endowed with partitions Π_A and Π_B , and let $E \subseteq \Omega$ be some given event.*

1. *If agents have the same prior probability P over Ω , and*
2. *if it is common knowledge at ω that $P(E | \Pi_A(\omega)) = p_A$ and $P(E | \Pi_B(\omega)) = p_B$,*

then $p_A = p_B$.

It is worth noting that stating "*Rational agents cannot agree to disagree*" is a slight abuse of language. It may be common knowledge among two agents that they hold different posterior probabilities. Let the set of states of the world be $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, and consider two agents A and B endowed with partitions $\Pi_A : \{\omega_1, \omega_2\}\{\omega_3, \omega_4\}$ and $\Pi_B :$

$\{\omega_1\}\{\omega_2, \omega_3\}\{\omega_4\}$. Suppose that they share a uniform prior P over Ω ($P(\omega_k) = 1/4 \forall k$), and that they express their posterior probability of the event $\{\omega_2, \omega_3\}$.¹

$$\Pi_A = \{\omega_1, \omega_2\}_{1/2}\{\omega_3, \omega_4\}_{1/2}$$

$$\Pi_B = \{\omega_1\}_0\{\omega_2, \omega_3\}_1\{\omega_4\}_0$$

Clearly, agents A and B hold different posteriors in every state of the world. As a consequence, it is effectively common knowledge in every state of the world among A and B that they have different posteriors: for all $\omega \in \Omega$, $M(\omega) = \Omega = \{\omega' \in \Omega \mid P(\{\omega_2, \omega_3\} \mid \Pi_A(\omega')) \neq P(\{\omega_2, \omega_3\} \mid \Pi_B(\omega'))\}$. Therefore, the result of Aumann [1976] does not state that “it cannot be common knowledge among rational agents that they disagree on their posteriors”, but that “if their posteriors are common knowledge, then they have to agree on their posteriors”.

The assumption of a common prior is central to Aumann’s result on the impossibility of agreeing to disagree, and is the basic assumption behind epistemic justifications of the concepts of correlated equilibrium (Aumann [1987]) and Nash equilibrium (Aumann and Brandenburger [1995]).² Criticisms of the common prior assumption concern its meaning in models of incomplete information.³ In those models, a state of the world describes individual belief hierarchies about the actual world, and is for Lipman [1995, p 2] “a fictitious construct, used to clarify our understanding of the real world”. Therefore, the assumption of a common prior on the set of states of the world “seem[s] to be based on giving the artificially constructed states more meaning than they have” (Dekel and Gul,

¹The subscripts describe individual posteriors in each cell.

²For a review of results about epistemic foundations of solution concepts in game theory, see Bonanno and Nehring [1997] for instance.

³Lipman [1995], Gul [1996].

[1997, p 115]). We do not want to discuss the plausibility (or the non-plausibility) of the common prior assumption in Aumann structures.⁴ However, we want to emphasize that the major contribution of Aumann's result does not depend on the common prior assumption, as we shall see in the next paragraph.

The proof of Aumann's result uses the following argument. Denote P_i the prior probability of agent i , and consider some event E . If $P_i(E | \Pi_i(\omega))$ is common knowledge at ω , then $P_i(E | \Pi_i(\omega))$ must be equal to $P_i(E | M(\omega))$. What does it mean? If i 's posterior probability is common knowledge at some state, then this posterior would have been the same, had i conditioned on the basis of the public information ${}^5M(\omega)$. Therefore, if each posterior $P_i(E | \Pi_i(\omega))$ is common knowledge at ω , then $P_i(E | \Pi_i(\omega)) = P_i(E | M(\omega))$ for all i . Imagine that each individual "permute" his information partition with the one of another individual. The partition of common knowledge does not change, and common knowledge of individual posteriors still implies that i 's posterior is equal to $P_i(E | M(\omega))$. As a consequence, common knowledge of individual posteriors implies that individual posterior probabilities do not reflect the differential information that each agent possesses. Note that this does not depend on the commonness of the prior at all. This implication of common knowledge is sometimes interpreted as the fact that *common knowledge negates asymmetric information*. As soon as individual posteriors are common knowledge, they do not depend on individuals' private information anymore, but only on the public information.

The Agreeing to Disagree literature raises basically the same question as Aumann's

⁴See Bonanno and Nehring [1999].

⁵Recall that public events are common knowledge whenever they occur. Therefore, they are cells or unions of cells of the partition of common knowledge.

result. To what extent differences in decisions can be explained on the basis of asymmetric information? Under what conditions common knowledge of individual decisions negates asymmetric information?

2.3. Common knowledge of individual decisions negates asymmetric information: the agreement theorems

Following Aumann [1976], a lot of papers have addressed the issue of how common knowledge of individual decisions might negate asymmetric information. These results are sometimes called the *agreement theorems*. In order to present these theorems, we use the following unified framework. We define a model m as a collection $m = (\Omega, (\Pi_i)_{i \in \mathcal{N}}, (\delta_i)_{i \in \mathcal{N}})$ where Ω is the set of states of the world, $(\Pi_i)_{i \in \mathcal{N}}$ the individual information partitions, and $(\delta_i)_{i \in \mathcal{N}}$ the individual decision rules. A decision rule $\delta_i : 2^\Omega \setminus \emptyset \rightarrow \mathcal{D}_i$ prescribes to agent i what decision to make as a function of any information situation i might be in. Posterior probabilities as in Aumann [1976], conditional expectations or discrete decisions such as “Buy” or “Sell” all correspond to particular decision rules. The decision made by agent i at state ω is generated by i 's decision rule δ_i , and is $\delta_i(\Pi_i(\omega))$. Let \mathcal{M} be the set of models. We define an outcome as a function $\phi : \mathcal{M} \times \Omega \rightarrow \Phi$, which associates to each vector $((\Omega, (\Pi_i)_{i \in \mathcal{N}}, (\delta_i)_{i \in \mathcal{N}}), \omega)$ a value in Φ .

Agreement theorems investigate under what conditions common knowledge of the outcome of the model implies that the outcome does not use in any way the differential information that each agent possesses. Namely, agreement theorems raise the following question:

What conditions are sufficient to guarantee that, for all model m and all outcome func-

tion ϕ satisfying these conditions, common knowledge of $\phi(m, \omega)$ at state ω implies that $\delta_i(\Pi_i(\omega)) = \delta_i(M(\omega))$ for all i ?

If such conditions are satisfied, and if, moreover, all individual decision rules are the same ($\delta_i = \delta \forall i$), then common knowledge of $\phi(m, \omega)$ at state ω implies that all agents must take the decision $\delta(M(\omega))$, a situation which is referred to as *consensus*. Most authors of agreement theorems studied the conditions under which common knowledge implies consensus, assuming commonness of the decision rules. We emphasize that the contribution of these results is not about consensus *per se*, but about the fact that common knowledge of individual decisions implies that all decisions are made on the basis of the same information.

We shall classify agreement theorems in two groups. Theorems in the first group investigate the conditions under which common knowledge of *the vector of individual decisions* negates asymmetric information. Theorems in the second group study what inferences can be made by individuals from common knowledge of *aggregate information* about individual decisions.

2.3.1. Common knowledge of individual actions

In this section, we consider that the outcome function ϕ is defined by

$$\phi((\Omega, (\Pi_i)_{i \in \mathcal{N}}, (\delta_i)_{i \in \mathcal{N}}), \omega) = ((\delta_i(\Pi_i(\omega)))_{i \in \mathcal{N}})$$

In this setting, Aumann [1976] proved that if δ_i is the posterior probability of an event $A \subseteq \Omega$, namely if $\delta_i(X) = P(A | X)$ for all $X \subseteq \Omega$, then common knowledge of $\delta_i(\Pi_i(\omega))$ for all i implies that $\delta_i(\Pi_i(\omega)) = \delta_i(M(\omega))$ for all i . As we illustrated it with the two traders example, the proof of this result uses the following property of *stability by disjoint*

union of posterior probabilities. Let $S = \{S_1, \dots, S_k\}$ be a family of disjoint events of Ω , and $A \subseteq \Omega$ some given event. If $P(A | S_l) = p \forall l = 1, \dots, k$, then $P(A | \bigcup_{l=1}^k S_l) = p$. Therefore, Aumann's result has been naturally extended to all decision rules satisfying this *stability by disjoint union* property, which was called *union consistency* by Cave [1983], and identified with the sure-thing principle by Bacharach [1985].

Definition 1 (Union consistency) *A function $f : 2^\Omega \rightarrow \mathcal{D}$ is union consistent if for all $E, E' \subseteq \Omega$ such $E \cap E' = \emptyset$, $f(E) = f(E') \Rightarrow f(E \cup E') = f(E) = f(E')$.*

Bacharach called this condition “sure thing principle” because its intuitive meaning sounds like Savage [1954]’s sure thing principle. If someone takes a particular decision whenever he knows that some event has occurred, and if he takes the same decision whenever he knows that this event has not occurred, then he need not be informed about the occurrence of this event to take this decision. For Bacharach [1985, p 168], it is a “certain fundamental principle of rational decision-making.” However, union consistency proved disputable, as it involves non-trivial assumptions whose appropriateness is questionable. We present Moses and Nachum [1990]’s criticism of union consistency in section 5.

Cave [1983] and Bacharach [1985] independently showed that if individual decision rules are union consistent, then agents cannot agree to disagree on their decisions.

Theorem 2 (Cave [1983], Bacharach [1985]) *Suppose that δ_i is union consistent for all i . If $\delta_i(\Pi_i(\omega))$ is common knowledge at state ω for all i , then $\delta_i(\Pi_i(\omega)) = \delta_i(M(\omega))$ for all i . If, moreover, $\delta_i = \delta \forall i$, then $\delta_i(\Pi_i(\omega)) = \delta_j(\Pi_j(\omega)) \forall j$.*

Posterior probabilities ($\delta(X) = P(E | X)$) and conditional expectations ($\delta(X) = E[Y | X]$) satisfy the union consistency property. Therefore, Cave and Bacharach’s result implies

some of the results in Milgrom and Stockey [1982] and in Rubinstein and Wolinski [1990]. Furthermore, Cave and Bacharach's result has direct implications for betting analysis. Consider the following particular decision rule. Let $X : \Omega \rightarrow \mathbb{R}$ be a real random variable, a a given real number, and the decision rule δ defined for all $F \subseteq \Omega$ by $\delta(F) = d_1 \Leftrightarrow E[X(\cdot) | F] \geq a$ and $\delta(F) = d_2 \Leftrightarrow E[X(\cdot) | F] < a$. Clearly, δ is union consistent. Therefore, Cave and Bacharach's result implies the one in Sebenius and Geanakoplos [1983]:

Theorem 3 (Sebenius and Geanakoplos [1983]) *Let $X : \Omega \rightarrow \mathbb{R}$ be a random variable, and a a real number. Consider two agents A and B endowed with partitions Π_A and Π_B of Ω . There is no state ω such that it is common knowledge to A and B at ω that $E[X | \Pi_A(\omega)] < a$ and $E[X | \Pi_B(\omega)] \geq a$.*

This result has several economic implications. First, it implies that two risk-neutral agents cannot bet against each other (it is therefore known as a *no-bet theorem*). Suppose that the variable X represents a bet between two risk-neutral agents. At state ω , agent A receives $X(\omega)$ and agent B receives $-X(\omega)$. If A and B accept the bet at state ω , then it becomes common knowledge among them at state ω that both expect a positive return from the bet. In other words, it is common knowledge at state ω that $E[X(\cdot) | \Pi_A(\omega)] > 0$ and that $E[-X(\cdot) | \Pi_B(\omega)] > 0$, which is $E[X(\cdot) | \Pi_B(\omega)] < 0$. This is impossible by Sebenius and Geanakoplos' result.

Another interesting application of this result deals with voting rules and group decision procedures. Consider the following example. On Friday night, Alice and Bob must decide whether they will go to the cinema (action C) or to the beach (action B) on Saturday. Their utility of both actions depends on Saturday's weather, which is sunny at state ω_1 , cloudy at state ω_2 and rainy at state ω_3 . If it is sunny, they both prefer to go to the beach

than to go to the cinema. If it is not, they both prefer to go to the cinema. Let Alice and Bob's common utility function be defined by $U(B, \omega_1) = 3$, $U(B, \omega_2) = U(B, \omega_3) = 0$ and $U(C, \omega) = 1$ for all ω . They share a uniform prior P over Ω , that is $P(\omega) = 1/3$ for all ω . They both receive a private signal about the weather, which leads them to have the following information partitions:

$$\Pi_A = \{\omega_1, \omega_2\}\{\omega_3\}$$

$$\Pi_B = \{\omega_1, \omega_3\}\{\omega_2\}$$

Alice's expected utility of going to the beach is $3/2$ at states ω_1 and ω_2 , and is 0 at state ω_3 . Her expected utility of going to the cinema is 1 in every state of the world. Therefore, if Alice had to decide by herself where to go, she would go to the beach at states ω_1 and ω_2 , and to the cinema at state ω_3 . If Bob had to decide by himself, he would go to the beach at states ω_1 and ω_3 , and to the cinema at state ω_2 . Suppose that state ω_1 occurs, namely that Saturday will be sunny. Alice does not know whether the true state is ω_1 or ω_2 . The action that maximizes her expected utility is going to the beach, but she knows there is half a chance for the weather to be cloudy and that she gets a payoff of 0. However, she knows that if the true state were ω_2 , namely if Saturdays' weather were cloudy, Bob would know it and would decide to go to the cinema. She also knows that if the true state is ω_1 , Bob would decide to go to the beach. As a consequence, when she phones Bob on Friday night she says to him "Let's do what you prefer." Bob makes the same reasoning. At state ω_1 , his optimal action is going to the beach, but he knows that he has a half chance of making a mistake. However, he also knows that Alice will be taking the correct decision in any state that he conceives as possible. Therefore, he also says to Alice "Let's do what you want." In this example, Alice and Bob both prefer that the other makes the decision for both of them at state 1. Note however that this fact is not common

knowledge among them, for at state ω_3 , they both want Alice to decide and at state ω_2 , they both want Bob to decide. More generally, let $(\Omega, (\Pi_i)_{i \in \mathcal{N}}, (\delta_i)_{i \in \mathcal{N}})$ be a model such that all agents have the same decision space ($\mathcal{D}_i = \mathcal{D} \forall i$). A voting rule is defined as a function $v : \mathcal{D}^{\mathcal{N}} \rightarrow \mathcal{D}$, which associates a particular decision to a decision profile. Majority rule, unanimity rule, Borda rule, dictatorship are some examples. At state ω , the decision profile is $\underline{d}(\omega) := (\delta_i(\Pi_i(\omega)))_{i \in \mathcal{N}}$. The result of the vote at state ω is then $v(\underline{d}(\omega))$. Suppose that individuals share a common utility function which depends both on the winning alternative, and on the state of the world: $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$. Agent i 's expected utility at state ω if voting rule v applies is $E[U(v(\underline{d}(\cdot)), \cdot) \mid \Pi_i(\omega)]$. An agent i prefers voting rule v to voting rule v' at ω if and only if $E[U(v(\underline{d}(\cdot)), \cdot) \mid \Pi_i(\omega)] \geq E[U(v'(\underline{d}(\cdot)), \cdot) \mid \Pi_i(\omega)] \Leftrightarrow E[(U(v(\underline{d}(\cdot)), \cdot) - U(v'(\underline{d}(\cdot)), \cdot)) \mid \Pi_i(\omega)] \geq 0$. As $U(v(\underline{d}(\cdot)), \cdot) - U(v'(\underline{d}(\cdot)), \cdot)$ is a particular random variable on Ω , Sebenius and Geanakoplos' result implies that it cannot be common knowledge in a group of agents that two of them disagree about the voting rule they should use to determine their collective decisions, provided that they have the same preferences.

The most well-known economic implication of the hypotheses that individual actions are common knowledge is the *no-trade theorem*. Milgrom and Stockey [1982] consider a pure exchange economy, where n risk averse agents have to trade in a situation of uncertainty about the state of the world. There are l commodities, and each agent's consumption set is \mathbb{R}_+^l . Each agent i is described by his initial endowment $e_i : \Omega \rightarrow \mathbb{R}_+^l$ (which is a random variable), his utility function $U_i : \Omega \times \mathbb{R}_+^l \rightarrow \mathbb{R}$ and his information partition Π_i . Agents share a common prior P over Ω . A trade t is a function from Ω into \mathbb{R}^{nl} , where $t_i(\omega)$ describes trader i 's net trade of commodities in state ω . A trade is feasible if each agent possesses a non-negative quantity of each good after trading, in every state of the world

$(e_i(\omega) + t_i(\omega) \geq 0 \forall i, \forall \omega)$, and if the sum of individual demands for some commodity in state ω is smaller than the total amount of this commodity in the economy in state ω ($\sum_{i=1}^n t_i(\omega) \leq 0 \forall \omega$). Milgrom and Stockey [1982] show that, if the initial allocation is Pareto-optimal, trade between rational, risk-averse agents, cannot be explained on the basis of asymmetric information.

Theorem 4 (Milgrom and Stockey [1982]) *Consider an economy where all traders are risk-averse, and where the initial allocation $(e_i(\omega))_i$ is Pareto-optimal in any state ω . If it is common knowledge at some state ω that t is a feasible trade, and that each trader weakly prefers $t(\omega)$ to the zero trade at ω , then every agent is indifferent between t and the zero trade at ω . If moreover all agents are strictly risk-averse, then $t(\omega)$ must be the zero trade.*

2.3.2. Common knowledge of an aggregate of individual decisions

It may sometimes be natural to assume that agents are facing aggregate information about others' decisions. Some agreement theorems study the aggregation of private information into a statistic, and the redistribution of information that occurs as individual make inferences from the common knowledge of this statistic. In Aumann [1976], agents have common knowledge of their posterior probabilities of some event. What happens if, for instance, they have common knowledge of the mean of their posteriors? In that case, agents cannot associate a particular posterior to a particular information partition when it comes to making inferences. McKelvey and Page [1986] show that if the statistic of individual posteriors satisfies a condition of *stochastic regularity*, then common knowledge of this statistic implies equality of individual posteriors. A stochastically regular function

is a one to one transformation of a *stochastically monotone* function, for which Bergin and Brandenburger [1990] provided a nice characterization. They showed that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is stochastically monotone if and only if f is additively separable into strictly increasing components, namely if it can be written in the form $f(x) = \sum_{i=1}^n f_i(x_i)$, where x_i denotes the i^{th} coordinate of x , and $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing for all i .

Theorem 5 (McKelvey and Page [1986]) *Consider n agents, each agent i being endowed with a prior P and an information partition Π_i of Ω , and let $E \subseteq \Omega$ be some given event. If ϕ is stochastically regular, then common knowledge of $\phi((P(E | \Pi_i(\omega)))_i)$ at state ω implies that $P(E | \Pi_i(\omega)) = P(E | M(\omega)) \forall i$.*

Nielsen, Brandenburger, Geanakoplos, McKelvey and Page [1990] extended McKelvey and Page's result from posterior probabilities to conditional expectations of some given random variable.

Theorem 6 (Nielsen *et al.* [1990]) *Consider n agents, each agent i being endowed with a prior P and an information partition Π_i of Ω , and let $X : \Omega \rightarrow \mathbb{R}$ be some random variable. If ϕ is stochastically regular, then common knowledge of $\phi((E[X(\cdot) | \Pi_i(\omega)])_i)$ at state ω implies that $E[X(\cdot) | \Pi_i(\omega)] = E[X(\cdot) | M(\omega)] \forall i$.*

Nielsen [1995] generalizes the last result to the case of random vectors. To do so, he adapts the definition of stochastic monotonicity to the multivariate case. In this setting, a function is stochastically monotone if it is additively separable into strictly comonotone components.

In Chapter 3 of this thesis, we will address the same question as McKelvey and Page's, in a more general setting where decisions may not be posterior probabilities.

Agreement theorems study the implications of common knowledge of a statistic of individual decisions, but not the conditions of emergence of such common knowledge situations. Consider Aumann [1976]'s setting. Obviously, without the assumption of common knowledge of individual posteriors, commonness of the priors does not guarantee equality of the posteriors. Consider the case where there are four states of the world, and where the two traders have a uniform prior over Ω ($P(\omega_k) = 1/4 \forall k$) and are endowed with the following partitions: $\Pi_A = \{\omega_1, \omega_2\}\{\omega_3, \omega_4\}$ and $\Pi_B = \{\omega_1, \omega_2, \omega_3\}\{\omega_4\}$. Let E be the event $\{\omega_2, \omega_3\}$. The subscript reflects individual posteriors of the event E :

$$\Pi_A = \{\omega_1, \omega_2\}_{1/2}\{\omega_3, \omega_4\}_{1/2}$$

$$\Pi_B = \{\omega_1, \omega_2, \omega_3\}_{2/3}\{\omega_4\}_0$$

At state ω_1 , A 's posterior probability of E is $1/2$, and B knows it, for it is the case in any state that B conceives as possible at state ω_1 . B 's posterior probability of E at state ω_1 is $2/3$, and A knows it, since it is the case in any state that A conceives as possible at state ω_1 . Therefore, A and B 's posteriors are *mutual knowledge*, but they differ because they are not *common knowledge*. In particular, B does not know that A knows that his posterior is $2/3$. At state ω_1 , B thinks that the true state of the world might be ω_3 . Therefore, B cannot exclude the possibility that A conceives ω_4 as possible, and thus believes that B 's posterior is 0.

The importance of the assumption that posteriors are common knowledge raises the question of how posterior probabilities, and more generally, how decisions, might become common knowledge. This question has no meaning in the state-space model of knowledge, where an event E being common knowledge is simply a property satisfied in some states and violated in others. However, how do we interpret the assertion that an event is common

knowledge when applying the model? It does not suffice for something to be common knowledge to simply tell everyone, since this ensures only that everyone knows, but not that everyone knows that everyone knows. It seems then difficult to observe a situation of common knowledge. However, if we impose sufficient structure on the interaction between agents, in particular if we assume that agents understand what they observe, and are able to make appropriate inferences from it, then we can model a communication process in which agents communicate their decisions until they become common knowledge to all.

2.4. Communication and common knowledge of individual decisions

In this section, we present some of the results which provided an answer to the second question addressed in the Agreeing to Disagree literature. Under what conditions communication of individual decisions can lead to common knowledge of decisions?

Let us consider the last example again. Suppose that A announces his posterior to B . In any state of the world, his posterior is $1/2$, hence B does not learn anything from A 's message. B 's partition remains:

$$\Pi_B : \{\omega_1, \omega_2, \omega_3\}_{2/3} \{\omega_4\}_0$$

Then B communicates his posterior to A . A knows that B will announce $2/3$ if she believes the state of the world to be in $\{\omega_1, \omega_2, \omega_3\}$, and will announce 0 otherwise. Therefore, B 's message drives A to distinguish states ω_1, ω_2 and ω_3 from state ω_4 . Since she could already distinguish states ω_1 and ω_2 from states ω_3 and ω_4 , his partition becomes:

$$\Pi_A : \{\omega_1, \omega_2\}_{1/2} \{\omega_3\}_1 \{\omega_4\}_0$$

A announces his posterior again. B is now able to distinguish states $\{\omega_1, \omega_2\}$ and state ω_3 .

His partition becomes:

$$\Pi_B : \{\omega_1, \omega_2\}_{1/2} \{\omega_3\}_1 \{\omega_4\}_0$$

From this time on, A and B do not learn any further information from communicating their posterior with each other. Why? Because their posteriors have become common knowledge to them. As a consequence, they are equal by Aumann's theorem.

This example illustrates the result of Geanakoplos and Polemarchakis [1982], who showed that two rational agents cannot “disagree forever”. Under the assumptions that information partitions are finite and agents have common priors, they showed that by communicating back and forth and revising their posteriors, the two agents will converge to a common posterior equilibrium, even though they may base their posterior on different information.

Theorem 7 (Geanakoplos and Polemarchakis [1982]) *Consider two agents endowed with finite partitions of Ω . If the two agents share a common prior, and if they alternately announce their posterior probability of a given event to one another, then their posterior probabilities will converge to a common posterior probability.*

This result was the first to provide an answer to the second question addressed in the literature, which deals with the conditions under which individual decisions might become common knowledge. Even if this is not the way the authors presented it, we could state Geanakoplos and Polemarchakis' result as follows. “If two agents alternately announce their posterior probabilities of a given event to one another, then eventually these posterior probabilities will become common knowledge to both agents. If, moreover, agents share

a common prior probability, then these posteriors must be equal by Aumann's theorem." Actually, Geanakoplos and Polemarchakis identified a particular communication protocol for two agents, according to which individual posteriors eventually become common knowledge between the two agents. We now present some of the results which investigated more generally how communication might generate common knowledge of individual decisions. We first properly define what is a communication protocol in the Agreeing to Disagree literature, and we present the way agents update their private information in such protocols.

2.4.1. *Communication and convergence of beliefs*

Following Geanakoplos and Polemarchakis [1982], some authors have analyzed the conditions under which communication might create common knowledge. Again, we shall use a unified framework to present the results on the topic. A model is a collection $(\Omega, (\Pi_i)_{i \in \mathcal{N}}, (\delta_i)_{i \in \mathcal{N}}, Pr)$, where Ω is the set of states of the world, $(\Pi_i)_{i \in \mathcal{N}}$ the individual partitions, $(\delta_i)_{i \in \mathcal{N}}$ the individual decision rules, and Pr the communication protocol. A communication protocol Pr is a pair of functions which determine the set of senders and the set of receivers of the communication at each date. Formally, $Pr = (s(\cdot), r(\cdot)) : \mathbb{N} \rightarrow 2^{\mathcal{N}} \times 2^{\mathcal{N}}$, where $s(t)$ and $r(t)$ respectively stand for the set of senders and of receivers of the communication which takes place at date t . At each date t , every sender $j \in s(t)$ communicates the private value of δ_j to every receiver $k \in r(t)$.

Communication is completely non strategic, namely if i 's private information⁶ is $X \subseteq \Omega$, then i truthfully communicates the value $\delta_i(X)$. Implicitly, one assumes that agents commit, or are constrained to communicate *via* decision rules δ_i .

⁶ X is the private information of an agent iff X is the smallest subset of Ω such that the agent thinks that the true state of the world belongs to X and not to $\Omega \setminus X$.

During the communication process, all receivers update their beliefs according to what they hear, and those who are not recipient of the communication at that date try to infer what recipients might have heard. We denote $\Pi_i(\omega, t)$ the set of possible states for an agent i at time t if the state of the world is ω . It is defined for all i and for all ω by the following recursive process:

$$\Pi_i(\omega, 0) = \Pi_i(\omega) \text{ and } \forall t \geq 1,$$

$$\Pi_i(\omega, t + 1) = \Pi_i(\omega, t) \cap \{\omega' \in \Omega \mid \delta_j(\Pi_j(\omega', t)) = \delta_j(\Pi_j(\omega, t)) \forall j \in s(t)\} \text{ if } i \in r(t),$$

$$\Pi_i(\omega, t + 1) = \Pi_i(\omega, t) \text{ otherwise.}$$

If $|\Omega| < \infty$ and $|\mathcal{N}| < \infty$, then there exists some $T < \infty$ such that $\forall t \geq T$, $\Pi_i(\omega, t) = \Pi_i(\omega, T)$ for all i . In the sequel, we will use the following notation:

- $\Pi_i^*(\omega)$ denotes the limiting value of $\Pi_i(\omega, t)$, namely $\Pi_i^*(\omega) := \lim_{t \rightarrow \infty} \Pi_i(\omega, t)$.
- Π_i^* denotes the information partition of agent i at the limit of the process, and will be called i 's limit information partition.
- M^* denotes the partition of common knowledge at the limit, namely $M^* = \bigwedge_{i \in \mathcal{N}} \Pi_i^*$, and will be called the limit partition of common knowledge.

Formally, the question raised by the papers on the topic is the following:

What conditions should be imposed on $(\delta_i)_{i \in \mathcal{N}}$ and on Pr to guarantee that communication eventually leads to common knowledge of individual decisions, and therefore that $\delta_i(\Pi_i^(\omega)) = \delta_i(M^*(\omega))$ for all ω ?*

One first has to assume that nobody is excluded from communication. This condition is satisfied by *fair* protocols. A protocol is fair if every participant in this protocol communicates directly or indirectly with every other participant, infinitely many times.

Definition 2 (Fair protocols) *A protocol Pr is fair if for all pair of players (i, j) , $i \neq j$, there exists an infinite number of finite sequences t_1, \dots, t_K , with $t_k \in \mathbb{N}$ for all $k \in \{1, \dots, K\}$, such that $i \in s(t_1)$ and $j \in r(t_K)$.*

It is easy to see why communication may not lead to common knowledge of individual decisions with a non-fair protocol. Consider two agents denoted A and B , endowed with a uniform prior probability over $\Omega = \{\omega_1, \omega_2, \omega_3\}$, and with the following partitions: $\Pi_A : \{\omega_1, \omega_2, \omega_3\}$, $\Pi_B : \{\omega_1, \omega_2\} \cup \{\omega_3\}$. Suppose that A is the only one allowed to communicate his posterior probability of event $\{\omega_1\}$. In every state of the world, A 's posterior is common knowledge to A and B , whereas B 's posterior is not. In particular, A does not know whether B 's posterior is 0 or $1/2$.

We must distinguish two kinds of protocols: *public* and *non-public* protocols.⁷ We define a public communication protocol as a protocol in which all agents are the recipient of the communication at any period.

Definition 3 (Public protocol) *A communication protocol is public if for all $t \in \mathbb{N}$, $r(t) = \mathcal{N}$.*

The reason for the distinction between public and non-public protocols is that the way by which communication leads to common knowledge of individual decisions essentially differs whether the protocol is public or not.

⁷We depart from the definition of Koessler [2001], according to which a protocol is public if there exists $t \in \mathbb{N}$ such that $r(t) = \mathcal{N}$.

2.4.2. Public communication protocols

In a fair and public protocol, communication of individual decisions leads to common knowledge of individual decisions, without any condition on the decision rules.

Proposition 1 *Let $|\mathcal{N}| < \infty$ and $|\Omega| < \infty$. If Pr is fair and public, then for all $(\delta_i)_i$, $M^*(\omega) \subseteq \{\omega' \in \Omega \mid \delta_i(\Pi_i^*(\omega')) = \delta_i(\Pi_i^*(\omega)) \forall i \in \mathcal{N}\}$ for all i and all ω . If, moreover, δ_i is union consistent, $\delta_i(\Pi_i^*(\omega)) = \delta_i(M^*(\omega)) \forall i \in \mathcal{N}$.*

Proof: Let us fix some state ω . If Pr is public, then $i \in r(t)$ for all i, t . Therefore, $\Pi_i(\omega, t)$ is defined for all i by:

$$\Pi_i(\omega, 0) = \Pi_i(\omega) \text{ and for all } t \geq 0$$

$$\Pi_i(\omega, t+1) = \Pi_i(\omega, t) \cap H(\omega, t) \text{ with } H(\omega, t) := \{\omega' \mid \delta_j(\Pi_j(\omega', t)) = \delta_j(\Pi_j(\omega, t)) \forall j \in s(t)\}.$$

Let T be such that for all $t \geq T$, for all i , $\Pi_i(\omega, t+1) = \Pi_i(\omega, t)$. Let $t \geq T$. By definition of $\Pi_i(\omega, t)$, $\Pi_i(\omega, t+1) = \Pi_i(\omega, t) \Rightarrow \Pi_i(\omega, t) \subseteq H(\omega, t)$ for all i , as Pr is public. Therefore, the smallest set of states containing $\Pi_i(\omega, t)$ for all i is also contained in $H(\omega, t)$, which implies that $H(\omega, t)$ is common knowledge at date t . It follows that for all $j \in s(t)$, $\delta_j(\Pi_j(\omega, t))$ is common knowledge at date t . As the protocol is fair, for all $i \in \mathcal{N}$, there exists some date $t \geq T$ such that $i \in s(t)$.

If moreover δ_i is union consistent, as $\Pi_i(\omega, t) = \Pi_i^*(\omega)$ for all i , and all ω , we have $\delta_j(\Pi_j^*(\omega)) = \delta_j(M^*(\omega))$ for all $j \in s(t)$. Therefore, for all $i \in \mathcal{N}$, $\delta_i(\Pi_i^*(\omega)) = \delta_i(M^*(\omega))$.

□

Therefore, if the protocol is fair and public, the sufficient conditions under which communication of individual decisions leads to a consensus in decisions are the same as in the case of common knowledge of individual decisions:

Proposition 2 (Cave [1983]) *If the communication protocol is fair and public, then*

- *equality of decision rules: $\delta_i = \delta_j \forall i, j$;*
- *union consistency of decision rules;*

are sufficient conditions to guarantee that eventually, all individual decisions are the same.

If the protocol is not public however, commonness and union consistency of individual decision rules are not sufficient for consensus to emerge in any fair protocol.⁸ Why? Because individual decisions may fail to ever become common knowledge in non-public protocols.

2.4.3. *Non-public protocols*

In non-public protocols, agents may privately communicate at some dates. The typical example is Parikh and Krasucki's *round-robin* protocol, where agents are sitting around a table and whisper their decisions to their left neighbor. Let us present an example where individual decisions fail to ever become common knowledge in a fair but non-public protocol. There are four states of the world, and three agents who communicate according to a round-robin protocol. Let us consider the very artificial decision rules defined as follows. Agent *A* decides *a* at states ω_1 and ω_2 , and *b* otherwise. Agents *B* and *C* take the decision *a* in any state. The three agents are endowed with the following partitions:

$$\Pi_A = \{\omega_1, \omega_2\}_a \{\omega_3, \omega_4\}_b$$

$$\Pi_B = \{\omega_1, \omega_2\}_a \{\omega_3, \omega_4\}_a$$

$$\Pi_C = \{\omega_1, \omega_2, \omega_3, \omega_4\}_a$$

A privately communicates his decision to *B*. The partition of common knowledge between *A* and *B* is $M^{AB} = \{\omega_1, \omega_2\} \{\omega_3, \omega_4\}$. As *A*'s decision is *a* in states ω_1 and ω_2 ,

⁸Parikh and Krasucki [1990, p 185] provide such an example, which we give in chapter 6.

and b in states ω_3 and ω_4 , A 's decision is common knowledge between A and B , and B does not learn anything from A . B 's decision is the same in every state of the world, then C does not learn anything from B , and similarly, A does not learn anything from C . Therefore, $\Pi_i^* = \Pi_i$ for $i = A, B, C$ and $M^* = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. B and C 's decisions are trivially common knowledge to all agents, but A 's are not, for A takes different decisions in ω_1 and ω_3 for instance.

Let us show that if the communication protocol is fair and public however, individual decisions must become common knowledge, as stated in Proposition 1. In any fair protocol, A will have to speak at some date. Then B will not learn anything but C will be able to distinguish states ω_1 and ω_2 from states ω_3 and ω_4 . From this time on, nobody will learn anything more, and we will have $\Pi_i^* = \Pi_i$ for $i = A, B$, but $\Pi_C^* = \{\omega_1, \omega_2\}\{\omega_3, \omega_4\}$. Therefore, $M^* = \{\omega_1, \omega_2\}\{\omega_3, \omega_4\}$, and A 's decision will effectively be common knowledge in every state.

This example shows that, contrary to the case of public protocols, one has to impose some conditions on decision rules for individual decisions to become common knowledge in non-public protocols. One way of achieving common knowledge of individual decisions by communicating is to create a consensus. Suppose that all agents follow the same decision rule, and that this decision rule guarantees that communication according to any fair and non-public protocols eventually leads to a consensus, namely a situation in which all agents take the same decision. Even in non-public protocols, once a consensus is reached, the consensus decision is, formally, common knowledge to all agents. Let $d \in \mathcal{D}$ be some decision, and let $Cons(d) = \{\omega \mid \delta(\Pi_i^*(\omega)) = d \forall i \in \mathcal{N}\}$ be the event that the agents have reached a consensus on decision d . Weyers [1992] showed that if $\omega \in Cons(d)$, then $M^*(\omega) \subseteq Cons(d)$. In other words, a consensus is reached in fair protocols if and only if

there is common knowledge of that consensus. As a consequence, if a consensus is reached at some state ω , then individual decisions must be common knowledge at that state, and every one must behave as if there were no asymmetric information:

$$\omega \in \text{Cons}(d) \Rightarrow M^*(\omega) \subseteq \{\omega' \mid \delta(\Pi_i^*(\omega')) = d \forall i \in \mathcal{N}\} \Rightarrow \delta(\Pi_i^*(\omega)) = \delta(M^*(\omega)) \forall i \in \mathcal{N}.$$

To sum up, communication of individual decisions in public protocols leads to common knowledge of individual decisions without any assumptions on decision rules (such as like-mindedness, union consistency, and so on), whereas communication according to non-public protocols leads to common knowledge of individual decisions if communication leads to consensus. In non-public protocols, common knowledge of individual decisions emerge *via* the consensus.

Therefore, to know what conditions guarantee that communication in non-public protocols leads to common knowledge of individual decisions, one has to know what conditions guarantee that communication leads to a consensus in non-public protocols. Parikh and Krasucki [1990] were the first to investigate that case. They show that if all individuals have the same decision rule, and if this decision rule is *convex*, which is a stronger requirement than union consistency, then communication eventually leads to consensus in any fair protocol.

Definition 4 (Convexity)

- A function $f : 2^\Omega \rightarrow \mathbb{R}$ is convex iff for all $E, E' \subseteq \Omega$ such that $E \cap E' = \emptyset$, there exists $\alpha \in]0, 1[$ such that $f(E \cup E') = \alpha f(E) + (1 - \alpha)f(E')$.
- A function $f : 2^\Omega \rightarrow \mathbb{R}$ is weakly convex iff for all $E, E' \subseteq \Omega$ such that $E \cap E' = \emptyset$, there exists $\alpha \in [0, 1]$ such that $f(E \cup E') = \alpha f(E) + (1 - \alpha)f(E')$.

Typically, posterior probabilities are convex functions. Weak and strong convexity, like union consistency, apply to disjoint unions of events. They therefore suffer the same shortcomings as union consistency with respect to their meaning when applied to sets of states of the world, as we will discuss in section 5.

Theorem 8 (Parikh and Krasucki [1990]) *If the communication protocol is fair, then*

- *equality of decision rules: $\delta_i = \delta$ for all i ;*
- *convexity of decision rules;*

are sufficient conditions to guarantee that eventually, all individual decisions are the same.

Parikh and Krasucki [1990] show that union consistency is not sufficient for consensus to emerge in any fair protocol for more than two agents. We present Parikh and Krasucki's example in chapter 6. They also show that weak convexity guarantees the consensus result for three agents, but not for more than three agents. Krasucki [1996] shows that Parikh and Krasucki's result holds for union consistent decision rules and more than two agents, provided that the protocol contains no cycle, which excludes typical communication networks as the circle.

Parikh and Krasucki's convexity condition may not apply in some contexts, as shown in the following example. An individual contemplates buying a car. Suppose that the set of available decisions is $\{\text{buy}, \text{not buy}\}$. Suppose that we re-label the decisions in \mathbb{R} , with 1 standing for *buy* and 0 standing for *not buy*. The convexity condition implies that if $\delta(X) = 0$ and $\delta(Y) = 1$ for some disjoint X, Y , then $\delta(X \cup Y) \in]0, 1[$, which does not correspond to any decision in $\{\text{buy}, \text{not buy}\}$. In Chapter 4, we give a new condition that should be applied to the common decision rule for consensus to emerge in any fair

communication protocol. This condition is that the decision rule must be the maximizer of a conditional expected utility. Contrary to convexity, this condition applies to any decision space.

Let us now make a point on whether commonness of the decision rules is *necessary* to have that communication leads to common knowledge of individual decisions. Parikh and Krasucki's result holds with different decision rules if all individual rules are a bijective transformation of the same convex function, namely if there exists a convex function f such that for all i , $\delta_i = h_i \circ f$ with h_i bijective. It is typically the case where some agents speak in Russian, some others speak in Spanish, and all agents understand only their own language and English. If the first agents have a Spanish-English dictionary, and the second agents a Russian-English dictionary, they can speak in Russian and Spanish and still converge to common knowledge. However, Parikh and Krasucki's result may not hold in the general case of different decision rules, as we show with the following example. Consider three agents denoted A , B and C who share a uniform prior over $\Omega = \{1, \dots, 8\}$, and who communicate in turn the value of different decision rules:

$$\delta_A(X) = P(\{2, 3\} | X), \delta_B(X) = P(\{3, 5\} | X), \delta_C(X) = P(\{2, 5\} | X)$$

The three agents are endowed with the partitions:

$$\Pi_A : \{1, 2\}_{1/2}\{3, 4\}_{1/2}\{5, 6\}_0\{7, 8\}_0$$

$$\Pi_B : \{1, 3\}_{1/2}\{2, 4\}_0\{5, 7\}_{1/2}\{6, 8\}_0$$

$$\Pi_C : \{1, 5\}_{1/2}\{2, 6\}_{1/2}\{3, 7\}_0\{4, 8\}_0$$

Consider the case where A speaks to B , who speaks to C , who speaks to A , and so on.

The partition of common knowledge among agents A and B is

$$M^{AB} : \{1, 2, 3, 4\}\{5, 6, 7, 8\}$$

As A 's decision is $1/2$ in states 1, 2, 3 and 4, and 2 in states 5, 6, 7 and 8, A 's decision is common knowledge among A and B at every state of the world. Therefore, agent B does not learn anything from A 's message. The partition of common knowledge among agents B and C is

$$M^{BC} : \{1, 3, 5, 7\}\{2, 4, 6, 8\}$$

The set of states where B takes the decision $1/2$ is $\{1, 3, 5, 7\}$, and the set of states where B decides 0 is $\{2, 4, 6, 8\}$. Again, agent C does not learn anything from B 's message. Finally, the partition of common knowledge among agents C and A is

$$M^{AC} : \{1, 2, 5, 6\}\{3, 4, 7, 8\}$$

The set of states where C decides $1/2$ is $\{1, 2, 5, 6\}$, and the set of states where C decides 0 is $\{3, 4, 7, 8\}$. As a consequence, agent A does not learn anything from C 's message.

However, individual decisions are not common knowledge to *all* agents. The partition of common knowledge in the group is $M = \{\Omega\}$, and each agent is taking a different decision in state 1 and in state 8.

In public protocols, common knowledge emerge independently of the emergence of a consensus, whereas in non-public protocols, common knowledge of individual decisions is implied by the consensus, and therefore can emerge only under particular conditions on individual decision rules. Since like-mindedness is necessary for consensus to obtain, it is also a necessary condition for common knowledge of individual decisions to obtain. This raises the question of the meaning of the usual like-mindedness assumption, which we discuss in the next section.

2.5. Some criticisms addressed to the Agreeing to Disagree literature

We identified three criticisms that have been addressed to the Agreeing to Disagree literature. The first one deals with the plausibility of the situation of common knowledge of individual decisions. We saw that one can build communication protocols in which individual decisions eventually become common knowledge. The last two were addressed by Moses and Nachum [1990], and deal with the union consistency condition and the assumption of commonness of decision rules.

2.5.1. Union consistency

Union consistency is the key assumption of agreement theorems. It is the necessary condition for common knowledge of individual decisions to imply that decisions do not reflect the differential information that each agent possesses. Even in Aumann's result, only the union consistency property of posterior probabilities is used to establish the result.

As stated by Cave and Bacharach, union consistency is a technical condition. However, Bacharach justified it by identifying union consistency with Savage's sure thing principle, arguing that union consistency characterizes rational individuals' decision rules. His interpretation of union consistency is the following. If I take the same decision whether I know that some event E has occurred or I know that $\neg E$ has occurred, then I still take the same decision if I am completely ignorant about the occurrence of E . Moses and Nachum [1990] pointed out an important flaw in Bacharach's interpretation of union consistency. Consider an agent i endowed with a partition Π_i and following a decision rule $\delta_i : 2^\Omega \rightarrow \mathcal{D}$. Consider two states ω and ω' such that $\Pi_i(\omega) \cap \Pi_i(\omega') = \emptyset$. If δ_i is union consistent, then

$\delta_i(\Pi_i(\omega)) = \delta_i(\Pi_i(\omega')) = d \Rightarrow \delta_i(\Pi_i(\omega) \cup \Pi_i(\omega')) = d$. Yet Moses and Nachum pointed out that “taking the union of states of knowledge in which an agent has differing knowledge does not result in a state of knowledge in which the agent is more ignorant. It simply does not result in a state of knowledge at all!” (Moses and Nachum [1990, p 156]).

We interpret Moses and Nachum’s criticism as follows. Union consistency could be identified with Savages’ sure thing principle in a one-decision maker setting, where states of the world are states of nature, for they describe only objective facts. Consider two states, ω_1 and ω_2 , such that it is raining in state ω_1 and not in state ω_2 . Consider a decision maker A who is able to distinguish between ω_1 and ω_2 . A ’ information partition is

$$\Pi_A = \{\omega_1\}\{\omega_2\}$$

In state ω_1 , A knows that it is raining, and in state ω_2 , A knows that it is not raining. Suppose now that the agent is not able to distinguish ω_1 and ω_2 , namely that his partition is

$$\Pi_A = \{\omega_1, \omega_2\}$$

Then in both states, A does not know that it is raining, and does not know that it is not raining. The event $\{\omega_1, \omega_2\}$ is effectively a state of knowledge in which A is more ignorant about the fact that it is raining.

Suppose now that there are two decision makers A and B , and there are endowed with the following partitions

$$\Pi_A = \{\omega_1\}\{\omega_2\}$$

$$\Pi_B = \{\omega_1, \omega_2\}$$

In both states, B knows that either A knows that it is raining, or A knows that it is not raining. Therefore, the event $\{\omega_1, \omega_2\}$ cannot be seen as an event where A is more

ignorant about the rain anymore, as it is the event “ B knows that either A knows that it is raining, or A knows that it is not raining”.

There are two levels in Moses and Nachum’s criticism. First, they argue that union consistency does not capture the intuitive meaning of Savage’s sure thing principle. Indeed, we saw that taking the union of information sets may not result in an information set where the agent is more ignorant in interactive settings. However, one need not identifying union consistency with the sure thing principle. Union consistency is a technical condition of stability by disjoint union, which is satisfied *de facto* by, for instance, posterior probabilities, argmax rules and conditional expectations. Obviously, decision rules which are only required to satisfy union consistency may be quite artificial. But union consistency is the less demanding requirement to guarantee that common knowledge of individual decisions negates asymmetric information.

The second level is that union consistency “requires the decision function to be defined in a manner satisfying certain consistency properties at what amount to impossible situations” (p 152). Indeed, we saw that it is quite artificial to impose that $\delta_A(\{\omega_1\}) = \delta_A(\{\omega_2\}) = d \Rightarrow \delta_A(\{\omega_1, \omega_2\}) = d$, as $\{\omega_1, \omega_2\}$ is not a possible information situation for A . However, we think it is worth noting that this level of Moses and Nachum’s criticism also applies to every result using union consistency, namely *every* result in the Agreeing to Disagree literature, in particular to Aumann’s agreement theorem. Aumann uses posterior probabilities. Yet if it makes no sense wondering what would be A ’s decision if he happened to know $\{\omega_1, \omega_2\}$, then it makes no sense either wondering what would be his probability of the event “It is raining” had he had the information $\{\omega_1, \omega_2\}$.

To conclude, we think that union consistency should only be seen as a technical requirement. However, decision rules which are only required to satisfy union consistency may be

quite artificial, and one should keep it in mind when applying results of the Agreeing to Disagree literature.

2.5.2. *Like-mindedness*

We saw in section 3 that commonness of individual decision rules (henceforth like-mindedness) is independent from the fact that common knowledge of individual decisions negates asymmetric information. However, it is a necessary condition for communication to lead to common knowledge of individual decisions in non-public protocols. It therefore raises the question of the relevance of like-mindedness assumptions.

Basically, the like-mindedness assumption states that given the same information, individuals must behave in the same way. This property has been formalized in the Agreeing to Disagree literature by the fact that agents follow the same decision rule. If quite natural, commonness of the decision rules implies some non-trivial hidden assumptions.

Consider a simple example, with three states of the world, and two agents endowed with the following partitions:

$$\Pi_A = \{\omega_1\}\{\omega_2, \omega_3\}$$

$$\Pi_B = \{\omega_1, \omega_2\}\{\omega_3\}$$

If A and B follow the same decision rule, then this decision rule must at least be defined on the same set of events (if not the entire set of events). Consider the event $E := \{\omega_1\}$. In state ω_1 , A knows E , and knows that B does not know E . Clearly, B cannot know that “ A knows E and A knows that B does not know E ”. Therefore, it makes no sense assuming that whenever B ’s private information is $\{\omega_1\}$, B makes the same decision as A when A ’s private information is $\{\omega_1\}$. In this example, B ’s decision rule should not be defined on $\{\omega_1\}$ and $\{\omega_2, \omega_3\}$, as these are not possible information sets for B .

The problem comes from the fact that decision rules are defined from the entire set of events into itself, associating decisions to set of states of the world that may not together form an information set. However, if $\{\omega_1\}$ is not an information set for B , it may well become one. For instance, if A communicates his private information to B , B 's information partition will become $\Pi'_B = \{\omega_1\}\{\omega_2\}\{\omega_3\}$, and the fact that δ_B is defined also on $\{\omega_1\}$ will make sense. To sum up, the basic problem with like mindedness is the same as the one behind the assumption that decision rules are defined on the whole set of events. However, it is justified in dynamic settings, where information sets evolve over time.

2.6. Conclusion

The Agreeing to Disagree literature addresses the following questions.

1. Under what conditions common knowledge of a statistic of individual decisions implies that decisions do not reflect the differential information that each individual possesses?
2. Under what conditions communication of individual decisions eventually leads to the situation of common knowledge of individual decisions?

As an answer to the first question, Cave [1983] and Bacharach [1985] showed that union consistency is a sufficient condition in the case where agents have common knowledge of their individual decisions. We saw that union consistency is the less demanding requirement which guarantees that common knowledge of individual decisions negates asymmetric information. McKelvey and Page [1986] showed that common knowledge of a stochastically regular statistic of individual posterior probabilities negates asymmetric information. In Chapter 3, we will provide another answer than McKelvey and Page's in the case where

individual decisions may not be posterior probabilities.

As an answer to the second question, Parikh and Krasucki [1990] showed that if individuals all follow the same convex decision rule, then communication eventually leads to consensus and common knowledge of individual decisions in any fair protocol. In public protocols, it is sufficient to assume that individual decision rules are union consistent, and commonness of the decision rules is not required. We saw that common knowledge of individual decisions emerge via the consensus in non-public protocols, and therefore requires commonness of the decision rules, whereas it emerges independently of any assumption on the decision rules in public protocols. In Chapter 4, we will provide a new condition on the common decision rule, sufficient for consensus and common knowledge of individual decisions to emerge in any fair protocol. Contrary to Parikh and Krasucki's convexity, our condition applies to any decision space.

This review of the Agreeing to Disagree literature is obviously not exhaustive. We made the choice to present only results to which our contributions could be related. In chapters 3, 4 and 5, we assume that knowledge is partitional. Therefore, we did not review those results which extended Aumann's result in a non partitional framework. However, it is worth mentioning that Geanakoplos [1989] and Samet [1990] showed that Aumann's result still holds when dropping the Negative Introspection axiom.

In chapters 4 and 5, we consider a dynamical setting in which agents learn by communicating with each other. We assume perfect communication, in the sense that messages always reach their receivers, and agents hold no uncertainty about that. Therefore, we only reviewed results investigating how *perfect* communication might create common knowledge of individual decisions. Heifetz [1996] and Koessler [2001] investigated the case of *noisy* communication. Heifetz showed that in this setting, a consensus may occur without being

common knowledge. In the same setting, Koessler showed that common knowledge fails to emerge in non-public and noisy protocols. He provided a general result for emergence of consensus without common knowledge for those protocols.

Chapter 3

Consensus and common knowledge of an aggregate of decisions

3.1. Introduction

Aumann [1976] showed that rational agents having common knowledge of their posterior probabilities of an event cannot disagree, in the sense that common knowledge of their posteriors implies the equality of these posteriors, provided that they have the same prior. Generalizing Aumann's theorem on the impossibility of agreeing to disagree, Cave [1983] and Bacharach [1985] showed that it is impossible for people following the same decision rule to take different decisions if these decisions are common knowledge, when the decision rule satisfies a *union consistency* condition.

It may sometimes be more natural to assume that agents have aggregate information about the others' decisions. On financial markets, stock traders don't know exactly which of them bought or sold the stock. All they know is the asset price that summarizes all moves on the market, that is, all individual decisions. McKelvey and Page [1986] investigate the effect that common knowledge of a statistic of the actions taken by a group of agents has on

individual actions, assuming that individual actions are individual posterior probabilities of some given event. They show that if the statistic satisfies a *stochastic regularity* condition, then common knowledge of it implies equality of all individual posterior probabilities.

In this chapter, we address the same question as McKelvey and Page's, in a more general setting where individual decisions may not be posterior probabilities. We consider a set of rational individuals, and we suppose that each individual's knowledge is described by a partition of the set of states of the world. All agents follow the same decision rule, which prescribes what decision to make as a function of any information situation they might be in. Each agent takes a decision based on his private information, and then a statistic of all decisions is made public. We investigate what conditions should be imposed on the decision rule and the statistic to guarantee that, in a group of individuals, common knowledge of the value of the statistic implies that everyone behaves as if there were no private information. In McKelvey and Page, these conditions are that 1) decisions are posterior probabilities of some event, and 2) the aggregate statistic is stochastically regular. We show that the same result holds for more general decision rules, provided a different condition on the statistic. Indeed, we show that if the decision rule is *balanced union consistent* and if the statistic is *exhaustive*, then individuals cannot take different decisions if this statistic is common knowledge, although beliefs might well remain different. *Balanced union consistency* is stronger than Cave [1983]'s union consistency, but is still obeyed by conditional probabilities, optimal actions and conditional expectations. The difference with union consistency is that it put some structure on the decisions made on the basis of non-disjoint events. The *exhaustiveness* condition imposes that the statistic should describe how many agents carry out each decision.

Even if the statistic is not initially common knowledge, the argument of Geanakop-

los and Polemarchakis [1982] can be adapted to show that an iterative announcement of the statistic eventually achieves the situation of common knowledge of the statistic. Furthermore, we give an example which shows that the consensus achieved by the common knowledge of the aggregate statistic might well be inefficient, in the sense that agents can agree on a decision different from the one they would have taken under perfect information. Hence it can happen that some agents took a better decision before knowing the public information than after.

The result of McKelvey and Page [1986] has been extended by Nielsen, Brandenburger, Geanakoplos, McKelvey and Page [1990] and by Nielsen [1995]. Nielsen *et al.* [1990] generalize the result from conditional probabilities of an event to conditional expectations of a random variable, and Nielsen [1995] from random variables to random vectors. Our contribution to the literature is to extend the results of McKelvey and Page [1986] and Nielsen *et al.* [1990] to arbitrary decision spaces (decisions can be posterior probabilities as well as ‘buy’ or ‘sell’), provided that the decision rule is *balanced union consistent*.

Section 2 describes the model. Section 3 defines *balanced union consistency* and *exhaustiveness* and develops the result. Section 4 concludes with a brief discussion of the relation between McKelvey and Page’s stochastic regularity condition and our exhaustiveness condition. Proofs are gathered in Section 5.

3.2. The model

Let Ω be the finite set of states of the world, and 2^Ω the set of possible events. There are N agents, each agent i being endowed with a partition Π_i of Ω . When the state $\omega \in \Omega$ occurs, agent i knows that the true state of the world belongs to $\Pi_i(\omega)$, which is the cell of i ’s partition that contains ω . We say that a partition Π is *finer* than a partition Π'

if and only if for all ω , $\Pi(\omega) \subseteq \Pi'(\omega)$ and there exists ω' such that $\Pi(\omega') \subset \Pi'(\omega')$. A partition Π' is coarser than a partition Π if and only if Π is finer than Π' . The partition Π_i represents the ability of agent i to distinguish between the states of the world. The coarser his partition is, the less precise his information is, in the sense that he distinguishes among fewer states of the world. As usual, we say that an agent i endowed with a partition Π_i knows the event E at state ω if and only if $\Pi_i(\omega) \subseteq E$. We define the meet of the partitions $\Pi_1, \Pi_2, \dots, \Pi_N$ as the finest common coarsening of these partitions, that is the finest partition M such that for all $\omega \in \Omega$ and for all $i = 1, \dots, N$, $\Pi_i(\omega) \subseteq M(\omega)$.

Common knowledge of an event E at some state ω is the situation that occurs when each agent knows E at ω , each agent knows that each of them knows E at ω , each agent knows that each agent knows that each agent knows... etc. Aumann [1976] showed that, given a set of N agents, the meet M of their N partitions is the partition of common knowledge among these N agents. Hence we say that an event E is common knowledge at state ω iff $M(\omega) \subseteq E$.

We suppose that each agent takes a decision d in a space \mathcal{D} . All agents follow the same decision rule, $\delta : 2^\Omega \setminus \emptyset \rightarrow \mathcal{D}$, which prescribes what decision to make as a function of any information situation they might be in. As the set of states of the world Ω is finite, the set of possible information situations is finite too. Consequently, the set of available decisions is also finite. Let $\delta(2^\Omega \setminus \emptyset) := \{d_1, \dots, d_m\}$ denote the range of the decision rule, with $m < \infty$.

Agents make their decisions on the basis of their private information. If the state ω occurs, agent i 's private information is $\Pi_i(\omega)$, therefore i takes the decision $\delta(\Pi_i(\omega))$. We note $\underline{\delta}_i(\omega)$ agent i 's decision at ω and $\underline{\delta}(\omega)$ the decision profile at ω . That is to say $\underline{\delta}_i(\omega) := \delta(\Pi_i(\omega))$ and $\underline{\delta}(\omega) := (\underline{\delta}_i(\omega))_i$. We investigate the effect that common knowledge

of a statistic Φ of all individual decisions at some state ω will have on $\underline{\delta}(\omega)$. We say that the statistic Φ is common knowledge at state ω if the event $\{\omega' \in \Omega \mid \Phi(\underline{\delta}(\omega')) = \Phi(\underline{\delta}(\omega))\}$ is common knowledge at ω .

3.3. The consensus result

As individuals act on the basis of their private information, the statistic of individual decisions is informative. We study what inferences can be made about $\underline{\delta}(\omega)$ from the public information $\Phi(\underline{\delta}(\omega))$. More precisely, our theorem investigates what conditions should be imposed on δ and Φ to guarantee that common knowledge of $\Phi(\underline{\delta}(\omega))$ among a set of individuals implies a consensus on their decisions.

The condition we impose on the decision rule is called *balanced union consistency*. Before stating it, let us define what we call a *balanced family* of 2^Ω , which slightly differs from Shapley [1967]'s definition.

Definition 1 *A non-empty family $\mathcal{B} \subseteq 2^\Omega$ is balanced if there exists a family of non-negative reals $\{\lambda_S\}_{S \in \mathcal{B}}$, called balancing coefficients, such that $\sum_{S \in \mathcal{B}, \omega \in S} \lambda_S = 1$ for every $\omega \in \bigcup_{S \in \mathcal{B}} S$.*

An example of a balanced family of $\Omega = \{1, 2, 3, 4, 5, 6\}$ is $\mathcal{B} = \{\{1, 2\}, \{3, 4\}, \{1, 2, 4\}, \{1, 2, 3\}\}$, which is balanced with respect to coefficients $\lambda_{\{1,2\}} = \lambda_{\{1,2,4\}} = \lambda_{\{1,2,3\}} = 1/3$ and $\lambda_{\{3,4\}} = 2/3$.

In Shapley [1967]'s definition, $\mathcal{B} \subseteq 2^\Omega$ is a balanced family if $\sum_{S \in \mathcal{B}, \omega \in S} \lambda_S = 1$ for every $\omega \in \Omega$ (whereas we require that is the case only for all $\omega \in \bigcup_{S \in \mathcal{B}} S$). Therefore, if a family \mathcal{B} is balanced according to Shapley's definition, then it is balanced according to ours.

We can now state the definition of balanced union consistency.

Definition 2 *The decision rule δ is balanced union consistent iff for every balanced family of events \mathcal{B} , $\delta(S) = d \forall S \in \mathcal{B} \Rightarrow \delta(\bigcup_{S \in \mathcal{B}} S) = d$.*

Balanced union consistency implies the standard union consistency condition of Cave [1983] and Bacharach [1985],¹ which imposes that if a decision maker makes the same action whether he knows E or F , where E and F are disjoint events, then he still makes the same action if he knows $E \cup F$. Union consistency tells nothing if E and F are not disjoint. We need the stronger condition of balanced union consistency to put some structure on the decisions made on the basis on non-disjoint events. However, balanced union consistency is still obeyed by usual decision rules in economics, in particular by *argmax rules*. We say that agents follow an argmax rule if they choose actions that maximize their expected utility given their private information. The argmax rule for an agent endowed with a utility function $U : \mathcal{D} \times \Omega$ and a prior belief P over Ω is defined from $2^\Omega \setminus \emptyset \rightarrow 2^\mathcal{D}$ by $\delta(X) = \operatorname{argmax}_{d \in \mathcal{D}} E[U(d, \cdot) | X] \subseteq \mathcal{D}$ for all $X \subseteq \Omega$.

Lemma 1 *Argmax rules are balanced union consistent.*

As posterior probabilities² and conditional expectations³ are particular argmax rules, they also obey balanced union consistency.

The sufficient condition we impose on the statistic is *exhaustiveness*. We say that a statistic is *exhaustive* if it is a one to one transformation of the statistic Φ^* defined as follows.

¹A family of disjoint events is always balanced.

²If $\mathcal{D} = [0, 1]$ and $U(d, \omega) = -1/2d^2 + d\mathbb{1}_A(\omega)$, then $\operatorname{argmax}_{d \in [0, 1]} E[U(d, \cdot) | X] = P(A | X)$.

³If $\mathcal{D} = \mathbb{R}$ and $U(d, \omega) = -1/2d^2 + dY(\omega)$, then $\operatorname{argmax}_{d \in \mathbb{R}} E[U(d, \cdot) | X] = E(Y | X)$.

Definition 3 $\Phi^* : \{d_1, \dots, d_m\}^N \rightarrow \mathbb{N}^m$ is defined by $\Phi^*(x_1, \dots, x_N) = (\sum_{i=1}^N \mathbb{1}_{x_i=d_1}, \dots, \sum_{i=1}^N \mathbb{1}_{x_i=d_m})$

In other words, the statistic Φ^* is a counting measure of individual decisions. A natural example is an opinion poll over the whole population. Suppose that people have to answer the following question: "Do you think the unemployment situation is: a) very worrying, b) pretty worrying, c) a little worrying, d) not worrying at all ". The decision is to choose one of the four alternatives. Suppose that there are ten agents, and that the sequence of their decisions is $(a, a, b, d, a, c, b, b, c, a)$. Then $\Phi^*((a, a, b, d, a, c, b, b, c, a)) = (4, 3, 2, 1)$.

Our theorem holds for any exhaustive statistic $\Phi = h \circ \Phi^*$, with h one to one. As our results do not depend on the transformation h at any point, we state the theorem and the proof for Φ^* .

Theorem 1 *If δ is balanced union consistent, then at every $\omega \in \Omega$, common knowledge of $\Phi^*(\underline{\delta}(\omega))$ implies that $\underline{\delta}_i(\omega) = \delta(M(\omega)) \forall i$.*

This theorem states that if all agents follow the same balanced union consistent decision rule, then common knowledge of the statistic Φ^* of individual decisions implies equality of all decisions. In other words, if the value of Φ^* is common knowledge at ω , then that value must be a permutation of the vector $(N, 0, \dots, 0)$. Let us emphasize that this result does not follow directly from the analysis of Cave [1983] and Bacharach [1985]. Basically, they show that it cannot be common knowledge that two *well identified* agents take different decisions. We show that it cannot be common knowledge that *there exist* two agents who take different decisions. The fact that decisions are anonymous in our model implies that agents have to base their inferences on less precise information than when they observe individual decisions.

Even if the aggregate $\Phi^*(\underline{\delta}(\omega))$ is not initially common knowledge at ω , we can adapt

the argument in Geanakoplos and Polemarchakis [1982] to show that repeated public announcements⁴ of the statistic must eventually lead to common knowledge of the statistic and hence, by Theorem 1, to equality of individual decisions.

We now discuss the roles played by our conditions in establishing the result. First, we assume that all agents follow the same decision rule ($\delta^i = \delta \forall i$) and that this is common knowledge among them. If individual decision rules $(\delta^i)_i$ are not identical, then common knowledge of $\Phi^*(\underline{\delta}(\omega))$ at state ω implies that, for all i , $\delta^i(\Pi_i(\omega')) = \delta^i(M(\omega))$ for all $\omega' \in M(\omega)$. In other words, the equality of individual decision rules is not necessary to obtain that “*common knowledge negates asymmetric information*”. However, it is obviously necessary to guarantee that consensus emerge. If $\delta^i \neq \delta^j$, it is then possible that $\delta^i(M(\omega)) \neq \delta^j(M(\omega))$, and thus that some agents take different decisions at ω .

We use exhaustiveness to prove that common knowledge of the statistic at some state ω implies common knowledge of a finite number of decision profiles, which are all permutations of the decision profile at ω . Obviously, it may not be the case if the statistic is not exhaustive, *e.g.* if the statistic is a constant function. Let us show that it may also not be the case for more interesting statistics than constant functions.

Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and consider three agents endowed with a utility function $U : \{0, 1/2, 1\} \times \Omega \rightarrow \mathbb{R}$ defined as follows:

$$U(x, \omega_1) = \begin{cases} 3 & \text{if } x = 1 \\ 0 & \text{if } x = 1/2 \\ 2 & \text{if } x = 0 \end{cases} \quad U(x, \omega_2) = \begin{cases} 0 & \text{if } x = 1 \\ 3 & \text{if } x = 1/2 \\ 2 & \text{if } x = 0 \end{cases} \quad U(x, \omega_3) = \begin{cases} 0 & \text{if } x = 1 \\ 3 & \text{if } x = 1/2 \\ 2 & \text{if } x = 0 \end{cases}$$

Agents follow the argmax rule associated with the utility function U and the uniform prior probability on Ω ($P(\omega) = 1/3 \forall \omega$). Suppose that the three agents are endowed with

⁴We show in chapter 6 that it may not be the case in settings of private communication, as in the one of Parikh and Krasucki [1990].

the following partitions:⁵

$$\Pi_1 : \{\omega_1, \omega_2\}_0 \{\omega_3\}_{1/2}$$

$$\Pi_2 : \{\omega_1, \omega_3\}_0 \{\omega_2\}_{1/2}$$

$$\Pi_3 = \{\omega_1\}_1 \{\omega_2, \omega_3\}_{1/2}$$

Imagine that the statistic of individual decisions is the mean, namely $\Phi(\underline{\delta}(\omega)) = \frac{1}{3} \sum_{i=1}^3 \delta_i(\omega)$ for all ω , then the value of the statistic is $1/3$ in every state of the world, although the decision profile is $(0, 0, 1)$ in state ω_1 , and $(0, 1/2, 1/2)$ in state ω_2 .

Balanced union consistency implies that permutations of the decision profile at some state cannot be common knowledge if agents take different decisions at this state. As the events corresponding to each decision profile may not be disjoint, we may not be able to apply Cave's union consistency. We now give an example showing that Cave's union consistency property does not imply consensus with an exhaustive statistic. Let $\Omega = \{1, 2, 3, 4, 5, 6\}$. The decision rule δ takes integer values and is defined as follows: let the subsets of Ω be numbered X_1, X_2, \dots , and let $n(X_k)$ be k . We let

$$\delta(\{1, 2, 3\}) = \delta(\{1, 4, 5\}) = \delta(\{2, 4, 6\}) = \delta(\{3, 5, 6\}) = 1$$

$$\delta(\{4, 5, 6\}) = \delta(\{2, 3, 6\}) = \delta(\{1, 3, 5\}) = \delta(\{1, 2, 4\}) = 2$$

and for all other $X \subseteq \Omega$, $\delta(X) = n(X) + 2$. This ensures that δ is union consistent. We consider four agents endowed with partitions Π_1, Π_2, Π_3 and Π_4 :

$$\Pi_1 : \{1, 2, 3\}_1 \{4, 5, 6\}_2$$

$$\Pi_2 : \{1, 4, 5\}_1 \{2, 3, 6\}_2$$

$$\Pi_3 : \{1, 3, 5\}_2 \{2, 4, 6\}_1$$

$$\Pi_4 : \{1, 2, 4\}_2 \{3, 5, 6\}_1$$

⁵The subscript reflects the decision taken in each cell.

At every state of the world, two agents take decision 1 and two agents take decision 2. Then it is common knowledge at every state of the world that the value of Φ^* is $(2, 2, 0, \dots, 0)$. Note however that this decision rule is not *balanced* union consistent. Denoting $\mathcal{B}_1 = \{\{1, 2, 3\}, \{1, 4, 5\}, \{2, 4, 6\}, \{3, 5, 6\}\}$, we have $\sum_{S \in \mathcal{B}_1, \omega \in S} 1/2 = 1$ for $\omega = 1, \dots, 6$. Denoting $\mathcal{B}_2 = \{\{4, 5, 6\}, \{2, 3, 6\}, \{1, 3, 5\}, \{1, 2, 4\}\}$, we have $\sum_{S \in \mathcal{B}_2, \omega \in S} 1/2 = 1$ for $\omega = 1, \dots, 6$. Hence \mathcal{B}_1 and \mathcal{B}_2 are balanced with respect to coefficients $\lambda_S = 1/2$ for every $S \in \mathcal{B}_1$ and every $S \in \mathcal{B}_2$. As $\delta(S) = 1$ for all $S \in \mathcal{B}_1$, if δ were balanced union consistent, we would have $\delta(\bigcup_{S \in \mathcal{B}_1} S) = 1$, that is to say $\delta(\{1, 2, 3, 4, 5, 6\}) = 1$. As $\delta(S) = 2$ for all $S \in \mathcal{B}_2$, if δ were balanced union consistent, we would have $\delta(\bigcup_{S \in \mathcal{B}_2} S) = 2$, that is to say $\delta(\{1, 2, 3, 4, 5, 6\}) = 2$, which brings the contradiction. However, for $N \leq 3$, union consistency of the decision rule is sufficient to guarantee the result.

3.4. Concluding remarks

We now briefly discuss the relative strength of the exhaustiveness condition relatively to the *stochastic regularity* of McKelvey and Page [1986]. A *stochastically regular* function $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is a one to one transformation of a *stochastically monotone* function. Bergin and Brandenburger [1990] showed that a function $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is stochastically monotone if and only if it can be written in the form $f(x) = \sum_{i=1}^N f_i(x_i)$ where each $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing.

Stochastically regular functions and exhaustive functions do not have the same domains. In order to compare these conditions, one first has to re-label the decision set \mathcal{D} in \mathbb{R} . Let us define a one to one function $t : \mathcal{D} \rightarrow \mathbb{R}$, and denote $\tilde{d}_k := t(d_k)$ for $k = 1, \dots, m$, and $\tilde{\mathcal{D}} := \{\tilde{d}_1, \dots, \tilde{d}_m\}$. Then one has to define what could mean exhaustiveness for real-valued functions, and what could mean stochastic monotonicity for non-real-valued functions. We

propose the following (obviously imperfect) definition. We say that:

- $f : \tilde{D}^n \rightarrow \mathbb{R}$ is “exhaustive” if $\exists h : \mathbb{R} \rightarrow \mathbb{N}^m$ such that $h \circ f(x) = \Phi^*(x)$.
- $g : \tilde{D}^n \rightarrow \mathbb{N}^m$ is “stochastically monotone” if $\exists h : \mathbb{N}^m \rightarrow \mathbb{R}$ such that $h \circ g(x) = \sum_{i=1}^n h_i(x_i)$, with $h_i(\cdot)$ strictly increasing for all i .

Clearly, some stochastically monotone functions are not exhaustive, in the sense we have just defined. The mean of individual posterior beliefs is clearly stochastically regular but may not be exhaustive. If for instance there are two agents and the value of f is 0.2, then both agents could have the posterior belief 0.2, or one could have the posterior 0.1 and the other the posterior 0.3.

Consider $\Phi^*(x) = (\sum_{i=1}^n \mathbb{1}_{x_i=\tilde{d}_1}, \dots, \sum_{i=1}^n \mathbb{1}_{x_i=\tilde{d}_m})$ and $h : \mathbb{N}^m \rightarrow \mathbb{R}$ defined by $h(z_1, \dots, z_m) := \tilde{d}_1 z_1 + \tilde{d}_2 z_2 + \dots + \tilde{d}_m z_m$. We have $h \circ \Phi^*(x) = \sum_{k=1}^m \tilde{d}_k \sum_{i=1}^n \mathbb{1}_{x_i=\tilde{d}_k} = \sum_{i=1}^n x_i$. Therefore, exhaustive functions are stochastically monotone according to the previous definition.

A very well-known fact in the Agreeing to Disagree literature is that consensus might well be inefficient, in the sense that the consensus decision may differ from the one that would have been made if all agents had shared their private information. The next example illustrates that point. Consider three agents who have to vote for candidate A or candidate B . The set of states of the world is $\{1, 2, 3, 4, 5\}$ and the three agents share a uniform prior P over Ω ($P(\omega) = 1/5 \forall \omega$). Agents have a utility function depending both on the elected candidate and the state of the world, defined by $U(A, \omega) = 1$ if $\omega \in \{2, 3\}$, $U(A, \omega) = 0$ if $\omega \in \{1, 4, 5\}$ and $U(B, \omega) = 1 - U(A, \omega)$ for all ω . Agents vote for the candidate that maximizes their expected utility, conditionally on their private information. In the particular setting we chose, this behavior can be described by the balanced union consistent rule $\delta : 2^\Omega \setminus \emptyset \rightarrow \{A, B\}$ defined by:

$$\left\{ \begin{array}{l} \delta(X) = A \Leftrightarrow P(\{2, 3\} | X) \geq 1/2 \\ \delta(X) = B \text{ otherwise} \end{array} \right.$$

This decision rule is an *argmax rule*, and is then balanced union consistent by Lemma

1. Agents are endowed with the partitions Π_1 , Π_2 and Π_3 .

$$\Pi_1 = \{1, 2\}_A, \{3\}_A, \{4, 5\}_B$$

$$\Pi_2 = \{1, 3\}_A, \{2, 4\}_A, \{5\}_B$$

$$\Pi_3 = \{1, 2, 3, 4, 5\}_B$$

The percentage of agents who intend to vote for candidate A is made public. This statistic is exhaustive as it allows to know how many agents intend to vote for each candidate. At states 1, 2 and 3, the percentage is $2/3$. At state 4, it is $1/3$, and it is 0 at state 5. Hence the announcement of the statistic leads agents to distinguish between states 1,2,3, state 4 and state 5. After the first announcement, individual partitions become:

$$\Pi_1 = \{1, 2\}_A, \{3\}_A, \{4\}_B \{5\}_B$$

$$\Pi_2 = \{1, 3\}_A, \{2\}_A \{4\}_B, \{5\}_B$$

$$\Pi_3 = \{1, 2, 3\}_A \{4\}_B \{5\}_B$$

The value of the statistic is now 1 at state 1,2,3, and 0 at state 4 and 5. As the partition of common knowledge is now $M = \{1, 2, 3\} \{4, 5\}$, the statistic has become common knowledge, and no agent can infer more information from it.

Note that in this example, the updating process achieves a decision consensus, but not a probability one. At state 1, one can notice that posterior probabilities for the event $\{2, 3\}$ of agents 1 and 2 are $1/2$, while the one of agent 3 is $2/3$. Furthermore, the consensus is inefficient at state 1, for the three agents would have known that the state of the world

was 1 had they shared their private information, and then would have elected candidate B instead of candidate A . Therefore, agent 3 took a better decision before receiving the public information than after, for he took the decision he would have taken if he had known the true state of the world.

3.5. Proofs

PROOF OF LEMMA 1:

Consider the argmax rule δ associated to a utility function U and a prior belief P , and let \mathcal{B} be a balanced family of events. As \mathcal{B} is balanced, there exist $\{\lambda_S\}_{S \in \mathcal{B}}$ such that $\lambda_S \geq 0 \forall S \in \mathcal{B}$ and $\sum_{S' \in \mathcal{B}, \omega \in S'} \lambda_{S'} = 1$ for every $\omega \in \bigcup_{S \in \mathcal{B}} S$. We have to show that if $\delta(S) = \delta(S')$ for all $S, S' \in \mathcal{B}$, then $\delta(\bigcup_{S' \in \mathcal{B}} S') = \delta(S)$ for all $S \in \mathcal{B}$. Consider $d \in \delta(S)$ and $d' \in \delta(\bigcup_{S' \in \mathcal{B}} S')$ and note $B = \bigcup_{S' \in \mathcal{B}} S'$.

- By definition of δ ,

$$E[U(d', \cdot) \mid B] \geq E[U(d, \cdot) \mid B]$$

- By definition of δ , for all $S \in \mathcal{B}$, $E[U(d, \cdot) \mid S] \geq E[U(d', \cdot) \mid S]$. It follows that for all S ,

$$\frac{1}{P(S)} \sum_{\omega \in \Omega} P(\omega) \mathbb{1}_{\omega \in S} U(d, \omega) \geq \frac{1}{P(S)} \sum_{\omega \in \Omega} P(\omega) \mathbb{1}_{\omega \in S} U(d', \omega)$$

which is equivalent to

$$\frac{1}{P(S)} \sum_{\omega \in B} P(\omega) \mathbb{1}_{\omega \in S} U(d, \omega) \geq \frac{1}{P(S)} \sum_{\omega \in B} P(\omega) \mathbb{1}_{\omega \in S} U(d', \omega)$$

As $\lambda_S \geq 0$, it follows that

$$\sum_{\omega \in B} P(\omega) \lambda_S \mathbb{1}_{\omega \in S} U(d, \omega) \geq \sum_{\omega \in B} P(\omega) \lambda_S \mathbb{1}_{\omega \in S} U(d', \omega)$$

Summing over S , we get:

$$\sum_S \sum_{\omega \in B} P(\omega) \lambda_S \mathbb{1}_{\omega \in S} U(d, \omega) \geq \sum_S \sum_{\omega \in B} P(\omega) \lambda_S \mathbb{1}_{\omega \in S} U(d', \omega)$$

$$\sum_{\omega \in B} P(\omega) U(d, \omega) \sum_S \lambda_S \mathbb{1}_{\omega \in S} \geq \sum_{\omega \in B} P(\omega) U(d', \omega) \sum_S \lambda_S \mathbb{1}_{\omega \in S}$$

Yet for every $\omega \in \Omega$, we have $\sum_S \lambda_S \mathbb{1}_{\omega \in S} = \sum_{S, \omega \in S} \lambda_S = 1$. Thus we have

$$\sum_{\omega \in B} P(\omega) U(d, \omega) \geq \sum_{\omega \in B} P(\omega) U(d', \omega)$$

which boils down to

$$E[U(d, \cdot) | B] \geq E[U(d', \cdot) | B]$$

We get $E[U(d, \cdot) | B] = E[U(d', \cdot) | B]$ for all $d \in \delta(S)$ and $d' \in \delta(B)$. Therefore, $\delta(\bigcup_{S \in \mathcal{B}} S) = \delta(S)$ for all S . \square

PROOF OF THEOREM 1:

Let ω be the state of the world. We note $K(\omega)$ the set of states of the world compatible with the value of the statistic $\Phi^*(\underline{\delta}(\omega))$:

$$K(\omega) = \{\omega' \in \Omega \mid \Phi^*(\underline{\delta}(\omega')) = \Phi^*(\underline{\delta}(\omega))\}$$

Given an agent i and a decision $d \in \mathcal{D}$, we note $K_i(d)$ the set of states (possibly empty) that are common knowledge at ω and in which i takes the decision d :

$$K_i(d) = M(\omega) \cap \{\omega' \in \Omega \mid \underline{\delta}_i(\omega') = d\}$$

For the rest of the proof, we consider a decision $d \in \mathcal{D}$ chosen by at least one agent at state ω , that is, such that $\exists i, \underline{\delta}_i(\omega) = d$, and we denote k the number of agents who take

the decision d at state ω :

$$k := \text{Card}(\{i \text{ s.t. } \underline{\delta}_i(\omega) = d\})$$

We first state three lemmas that will be used in the proof of the theorem.

Lemma 2 *If $M(\omega) \subseteq K(\omega)$, then $\{K_1(d), \dots, K_N(d)\}$ is a balanced family of $M(\omega)$.*

Proof: Let \mathcal{B} denote $\{K_1(d), \dots, K_N(d)\}$. For all i , $K_i(d) \subseteq M(\omega)$, and \mathcal{B} is non-empty as

$\exists i$ such that $\omega \in K_i(d)$. By definition of Φ^* , the fact that exactly k individuals take the decision d at ω implies that $\forall \omega' \in K(\omega)$, exactly k individuals take the decision d at ω' . As a consequence, if $M(\omega) \subseteq K(\omega)$, then $\forall \omega' \in M(\omega)$, there are exactly k agents who take decision d at state ω' . Then $\forall \omega' \in M(\omega)$, $\sum_{i=1}^N \mathbb{1}_{\omega' \in K_i(d)} = k$.

Denoting $\lambda_S = 1/k$ for all $S \in \mathcal{B}$, we have $\forall \omega' \in M(\omega)$, $\sum_{S \in \mathcal{B}, \omega' \in S} \lambda_S = 1$. \square

Lemma 3 *If δ is balanced union consistent, then $\delta(K_i(d)) = d$ for all i such that $K_i(d) \neq \emptyset$.*

Proof: If $K_i(d) \neq \emptyset$, then $K_i(d)$ is a union of cells of Π_i such that $\delta(\Pi_i(k)) = d$. If δ is

balanced union consistent, δ is also union consistent and $\delta(K_i(d)) = d$. \square

Lemma 4 *If $M(\omega) \subseteq K(\omega)$, then $\bigcup_{i=1}^N K_i(d) = M(\omega)$.*

Proof: By definition, $K_i(d) \subseteq M(\omega)$ for all i , then $\bigcup_{i=1}^N K_i(d) \subseteq M(\omega)$. If $M(\omega) \subseteq K(\omega)$,

then for all $\omega' \in M(\omega)$, $\exists i$ such that $\omega' \in K_i(d)$. Then $M(\omega) \subseteq \bigcup_{i=1}^N K_i(d)$. \square

We now turn to the proof of the theorem itself. If $\Phi^*(\underline{\delta}(\omega))$ is common knowledge at ω , then $M(\omega) \subseteq K(\omega)$. By lemma 2, $\{K_1(d), \dots, K_N(d)\}$ is balanced. If δ is balanced union consistent, by lemma 3 we have $\delta(K_i(d)) = d$ for all i such that $K_i(d) \neq \emptyset$, and then $\delta(\bigcup_{i=1}^N K_i(d)) = d$. Yet by lemma 4, $\bigcup_{i=1}^N K_i(d) = M(\omega)$. As a consequence, $\delta(M(\omega)) =$

d. As it is the case for any decision d taken by at least one agent at state ω , we have

$$\underline{\delta}_i(\omega) = \delta(M(\omega)) \forall i. \quad \square$$

Consensus, communication and knowledge: an extension with Bayesian agents

4.1. Introduction

Geanakoplos and Polemarchakis [1982] were the first to extend Aumann's result on the impossibility of agreeing to disagree to a dynamic framework. They showed that communication of posterior probabilities leads to a situation of common knowledge of these posteriors. Cave [1983] and Bacharach [1985] extended this result considering *union consistent*¹ functions more general than posterior probabilities. In all of these settings, communication is public, as achieved *e.g.* by auctions. Parikh and Krasucki [1990] investigated the case where communication is not public but in pairs. They defined an updating process along which agents communicate with each other, according to a protocol upon which they have agreed beforehand. At each stage one of the agents transmits to another agent the value of a certain function f , which depends on the set of states of the world he conceives as

¹Let Ω be the set of states of the world. $f : 2^\Omega \rightarrow \mathcal{D}$ is union consistent if $\forall X, Y \in 2^\Omega$ such that $X \cap Y = \emptyset$, $f(X) = f(Y) \Rightarrow f(X \cup Y) = f(X) = f(Y)$.

possible at that stage. Parikh and Krasucki [1990] showed that two conditions guarantee that eventually, all agents will communicate the same value (a situation we will refer to as a consensus): 1) a *fairness* condition on the communication protocol, which imposes that every agent has to be sender and receiver of the communication infinitely many times; 2) a *convexity* condition on the function whose value is communicated. Let Ω be the set of states of the world. A function $f : 2^\Omega \rightarrow \mathbb{R}$ is convex if $\forall X, Y \in 2^\Omega$ such that $X \cap Y = \emptyset$, there exists $\alpha \in]0, 1[$ such that $f(X \cup Y) = \alpha f(X) + (1 - \alpha)f(Y)$. This condition is satisfied by conditional probabilities for instance, and is more restrictive than Cave's union consistency.

Parikh and Krasucki's convexity condition may not apply in some contexts, as shown in the following example. An individual contemplates buying a car. The set of available decisions is $\{buy, not\ buy\}$. Suppose that we re-label the decisions in \mathbb{R} , with for instance 1 standing for *buy* and 0 standing for *not buy*. The convexity condition implies that if $f(X) = 0$ and $f(Y) = 1$ for some X, Y such that $X \cap Y = \emptyset$, then $f(X \cup Y) \in]0, 1[$, which does not correspond to any decision in $\{buy, not\ buy\}$. Hence there are some decision spaces for which, even after a re-labelling in \mathbb{R} , we may not be able to apply the convexity condition.

Parikh and Krasucki [1990] showed by a counter-example that *weak convexity*² and union consistency are not sufficient to guarantee that consensus occurs in any fair protocol. Krasucki [1996] investigated what restrictions on the communication protocol should be imposed to guarantee the consensus with any union consistent function. He showed that if the protocol is fair and contains no cycle, then communication of the value of any union

²Let Ω be the set of states of the world. $f : 2^\Omega \rightarrow \mathbb{R}$ is weakly convex if $\forall X, Y \in 2^\Omega$ such that $X \cap Y = \emptyset$, there exists $\alpha \in [0, 1]$ such that $f(X \cup Y) = \alpha f(X) + (1 - \alpha)f(Y)$.

consistent function leads to consensus.

In this chapter, we give a new condition on f for consensus to emerge in any fair communication protocol. This condition is that the function whose values are communicated is the maximizer of a conditional expected utility. Contrary to Parikh and Krasucki's convexity condition, this condition applies to any action space.

Even after an appropriate re-labelling of the image of f in \mathbb{R} , the functions we consider may not be representable by weakly convex functions. Furthermore, there exist weakly convex functions that do not obey our condition. Hence the class of functions we look at have a non-empty intersection with the class of weakly convex functions, but there is no inclusion relation among them. On the other hand, for any decision space, the functions we consider are union consistent.

4.2. Reaching a consensus

Let Ω be a finite set of states of the world. We consider a group of N agents, each of them endowed with a partition Π_i of Ω . All agents share some prior belief P on Ω . We note $\Pi_i(\omega)$ the cell of Π_i that contains ω . $\Pi_i(\omega)$ is the set of states that i judges possible when state ω occurs. As in Parikh and Krasucki [1990], agents communicate the value of a function $f : 2^\Omega \rightarrow \mathcal{D}$, according to a *fair* protocol Pr . A protocol is a pair of functions $(s(\cdot), r(\cdot)) : \mathbb{N} \rightarrow \{1, \dots, N\}^2$ where $s(t)$ stands for the sender and $r(t)$ the receiver of the communication which takes place at time t . A protocol is *fair* if no participant is blocked from the communication, that is if every agent is a sender and a receiver infinitely many times, and everyone receives information from every other, possibly indirectly, infinitely many times. Formally, given a protocol $(s(t), r(t))$, consider the directed graph whose

vertices are the participants $\{1, \dots, N\}$ and such that there is an edge from i to j iff there are infinitely many t such that $s(t) = i$ and $r(t) = j$. The protocol is fair if the graph above is strongly connected.

Except fairness, we do not make any assumption on the protocol. We assume that \mathcal{D} can be any compact subset of a topological space.

Agents share a common payoff function $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$, which depends on the chosen action $d \in \mathcal{D}$ and on the realized state of the world. We assume that $U(\cdot, \omega)$ is continuous on \mathcal{D} for all ω . What is communicated by an agent is the action that maximizes his expected utility, computed with respect to the common belief P .

In order to avoid indifference cases, we make the assumption that given any event, all actions have different expected utility conditional on this event.

Assumption [No indifference] *For any event $F \subseteq \Omega$, $\forall d \neq d' \in \mathcal{D}$, $E(U(d, \cdot) | F) \neq E(U(d', \cdot) | F)$.*

Without the no indifference assumption, the set of maximizing actions of an agent may not be a singleton, and we would have to specify the way agents choose between indifferent actions. It is clearly a strong assumption and we discuss the implications of relaxing it in the conclusion.

The function $f : 2^\Omega \rightarrow \mathcal{D}$ is then defined by:

$$\forall E \subseteq \Omega, f(E) = \operatorname{argmax}_{d \in \mathcal{D}} E(U(d, \cdot) | E)$$

Suppose now that Pr is some given protocol. The set of possible states for an agent i at time t if the state of the world is ω is denoted $C_i(\omega, t)$ and is defined by the following

recursive process:

$$C_i(\omega, 0) = \Pi_i(\omega)$$

$$C_i(\omega, t + 1) = C_i(\omega, t) \cap \{\omega' \in \Omega \mid f(C_{s(t)}(\omega', t)) = f(C_{s(t)}(\omega, t))\} \text{ if } i = r(t),$$

$$C_i(\omega, t + 1) = C_i(\omega, t) \text{ otherwise.}$$

The next result states that for all ω , $f(C_i(\omega, t))$ has a limiting value which does not depend on i .

Theorem 1 *There is a $T \in \mathbb{N}$ such that for all ω , i , and all $t, t' \geq T$, $C_i(\omega, t) = C_i(\omega, t')$.*

Moreover, if the protocol is fair, then for all i, j , for all ω , $f(C_i(\omega, T)) = f(C_j(\omega, T))$.

We discuss the properties of the function f defined as the argmax of an expected utility in chapter 6. First, f is clearly *union consistent* for any action space. Second, f may not be representable by a weakly convex function, namely a one to one function $g : \mathcal{D} \rightarrow \mathbb{R}$ may fail to exist such that $g \circ f$ is weakly convex. If such a function g were to exist, learning and consensus properties of f and $g \circ f$ would be the same. Therefore, the functions f we consider would be particular weakly convex functions, for which consensus obtains in any fair protocol. We show that it is not the case in chapter 6. Finally, there exist weakly convex functions that cannot be defined as the argmax of an expected utility. Such an example can be found in Parikh and Krasucki [1990, p 185]: they exhibit a weakly convex function f such that consensus may fail to occur in some protocols. We show in chapter 6 that it is not possible to find a utility function U and a probability P such that this function f is the argmax of the conditional expectation of U .

4.3. Concluding remarks

The assumption that for any subset F and any $d \neq d' \in \mathcal{D}$, $E[U(d, \cdot) | F] \neq E[U(d', \cdot) | F]$ is useful to avoid indifference cases, but is quite strong. We now discuss the implications of relaxing this assumption for the consensus result. The primary consequence is obviously that given a subset $F \subseteq \Omega$, $f(F)$ may not be a singleton. The fact that $f(\cdot)$ may not be a singleton does not matter for the result, as long as the function whose values are communicated is union consistent. More precisely, as long as the Result 1 of the proof holds, what action is communicated in the argmax does not matter, as they are all maximizing actions. Two cases can be identified.

1) First, agents could communicate the entire set of their maximizing actions. In that case, the result of Theorem 1 still holds, for f is still union consistent.

Proposition 1 *For all utility function $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$ and all probability P on Ω , the function $f : 2^\Omega \rightarrow 2^\mathcal{D}$ defined by $f(X) = \operatorname{argmax}_{d \in \mathcal{D}} E_P[U(d, \cdot) | X]$ for all $X \subseteq \Omega$ is union consistent.*

As f is union consistent, Result 1 of the proof holds. The rest of the proof is then the same as when the argmax is a singleton, but is more tedious.

2) Second, agents could be able to communicate only one action (which is more realistic). In that case, the emergence of a consensus depends on the way agents choose the action to be communicated in their argmax set. Let $f^*(E)$ denotes the action effectively communicated by an agent with private information E .

If they follow no particular selection rule, then the result of Theorem 1 does not hold, for the function f^* is not union consistent anymore. However, it is possible to find selection rules such that the result still holds if agents communicate one of their maximizing actions

according to these rules. Suppose that the decision space is endowed with a pre-order \succ , and that agents communicate their larger (or their smaller) maximizing action according to \succ . The so-defined function f^* is union consistent.

Proposition 2 *The function $f^* : 2^\Omega \rightarrow \mathcal{D}$ defined for any $F \subseteq \Omega$ by $f^*(F) = \operatorname{argmax}_{d \in f(F)} \succ$ ³ is union consistent.*

Note however that assuming such a selection rule is tantamount to making a no-indifference assumption.

4.4. Proofs

Proof: [Theorem 1]

1) As Ω is finite, the first part of the theorem is evident. In the sequel, we will note $C_i(\omega)$ the limiting value of $C_i(\omega, t)$, and C_i the information partition of agent i at equilibrium.

2) As in Parikh and Krasucki [1990], we prove the second part of the theorem for $N = 3$ and for a “round-robin protocol”, namely such that for all t , $s(t) = t \bmod 3$ and $r(t) = (t + 1) \bmod 3$. Note that this is sufficient to prove the theorem for any fair protocol. Our argument only uses the fact that we are able to find a chain $t_1 < t_2 < \dots < t_p$, with $T \leq t_1$, such that: (a) $s(t_1) = 1$, (b) the receiver at t_j is the sender at t_{j+1} , (c) the chain passes through all participants, finally returning to 1. This is implied by the fact that the protocol is fair.

Let M_{ij} be the partition of common knowledge among agents i and j at equilibrium, that is M_{ij} is the finest partition of Ω such that $\forall \omega$, $C_i(\omega) \subseteq M_{ij}(\omega)$ and $C_j(\omega) \subseteq M_{ij}(\omega)$. By consequence, $\forall \omega$, $M_{ij}(\omega)$ is a disjoint union of cells of C_i and a disjoint union of cells of C_j . $\sum_{C_i(k) \subseteq M_{ij}(\omega)}$ will denote the sum on all cells of C_i composing $M_{ij}(\omega)$.

³Or $f^*(F) = \operatorname{argmin}_{d \in f(F)} \succ$

At equilibrium, agent 1 communicates his optimal action to agent 2, agent 2 communicates his optimal action to agent 3 and agent 3 communicates his optimal action to agent 1. By consequence, the action taken by agent 1 is common knowledge among 1 and 2. Hence we have for all ω :

$$M_{12}(\omega) \subseteq \{\omega' \in \Omega \mid f(C_1(\omega')) = f(C_1(\omega))\}$$

As $M_{12}(\omega)$ is a disjoint union of cells of C_1 , union consistency of f implies that $f(M_{12}(\omega)) = f(C_1(k)) \forall k \in M_{12}(\omega)$.

• **Result 1** $E(U(f(M_{12}(\omega)), \cdot) \mid M_{12}(\omega)) = E[E(U(f(C_1(\cdot)), \cdot) \mid C_1(\cdot)) \mid M_{12}(\omega)]$

Proof: For all $\omega' \in M_{12}(\omega)$, $f(C_1(\omega')) = f(M_{12}(\omega))$. Then $E[E(U(f(C_1(\cdot)), \cdot) \mid C_1(\cdot)) \mid M_{12}(\omega)] = E[E(U(f(M_{12}(\omega)), \cdot) \mid C_1(\cdot)) \mid M_{12}(\omega)]$. As M_{12} is coarser than C_1 , the law of iterated expectations implies that $E[E(U(f(M_{12}(\omega)), \cdot) \mid C_1(\cdot)) \mid M_{12}(\omega)] = E(U(f(M_{12}(\omega)), \cdot) \mid M_{12}(\omega))$.

• **Result 2** $E(U(f(M_{12}(\omega)), \cdot) \mid M_{12}(\omega)) \leq \sum_{C_2(k) \subseteq M_{12}(\omega)} \frac{P(C_2(k))}{P(M_{12}(\omega))} E(U(f(C_2(k)), \cdot) \mid C_2(k))$

Proof: By definition, $\forall k \in M_{12}(\omega)$ we have:

$$E(U(f(M_{12}(\omega)), \cdot) \mid C_2(k)) \leq E(U(f(C_2(k)), \cdot) \mid C_2(k))$$

It implies that:

$$\sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(M_{12}(\omega)), \cdot) \mid C_2(k)) \leq \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) \mid C_2(k))$$

that is:

$$P(M_{12}(\omega)) E(U(f(M_{12}(\omega)), \cdot) \mid M_{12}(\omega)) \leq \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) \mid C_2(k)) \square$$

• **Result 3** $\forall i, j, E[E(U(f(C_i(\cdot)), \cdot) | C_i(\cdot))] = E[E(U(f(C_j(\cdot)), \cdot) | C_j(\cdot))]$

Proof:

$$E[E(U(f(C_1(\cdot)), \cdot) | C_1(\cdot))] = \sum_{M_{12}(\omega) \subseteq \Omega} P(M_{12}(\omega)) E[E(U(f(C_1(\cdot)), \cdot) | C_1(\cdot)) | M_{12}(\omega)]$$

Yet by results **1** and **2**, we have

$$P(M_{12}(\omega)) E[E(U(f(C_1(\cdot)), \cdot) | C_1(\cdot)) | M_{12}(\omega)] \leq \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) | C_2(k))$$

Then

$$\begin{aligned} E[E(U(f(C_1(\cdot)), \cdot) | C_1(\cdot))] &\leq \sum_{M_{12}(\omega) \subseteq \Omega} \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) | C_2(k)) \\ &= \sum_{C_2(k) \subseteq \Omega} P(C_2(k)) E(U(f(C_2(k)), \cdot) | C_2(k)) \\ &= E[E(U(f(C_2(\cdot)), \cdot) | C_2(\cdot))] \end{aligned}$$

Applying the same reasoning, we get

$$E[E(U(f(C_2(\cdot)), \cdot) | C_2(\cdot))] \leq E[E(U(f(C_3(\cdot)), \cdot) | C_3(\cdot))]$$

and

$$E[E(U(f(C_3(\cdot)), \cdot) | C_3(\cdot))] \leq E[E(U(f(C_1(\cdot)), \cdot) | C_1(\cdot))]$$

Hence $E[E(U(f(C_i(\cdot)), \cdot) | C_i(\cdot))] = E[E(U(f(C_j(\cdot)), \cdot) | C_j(\cdot))]$ for all i, j . \square

• **Result 4** For all $\omega \in \Omega$, we have

$$E(U(f(C_1(\omega)), \cdot) | C_2(\omega)) = E(U(f(C_2(\omega)), \cdot) | C_2(\omega))$$

$$E(U(f(C_2(\omega)), \cdot) | C_3(\omega)) = E(U(f(C_3(\omega)), \cdot) | C_3(\omega))$$

$$E(U(f(C_3(\omega)), \cdot) | C_1(\omega)) = E(U(f(C_1(\omega)), \cdot) | C_1(\omega))$$

Proof:

By **Result 3**, the inequality can not be strict in **Result 2**. Then we have:

$$P(M_{12}(\omega)) E(U(f(M_{12}(\omega)), \cdot) | M_{12}(\omega)) = \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) | C_2(k))$$

By definition, $E(U(f(C_1(k)), \cdot) | C_2(k)) \leq E(U(f(C_2(k)), \cdot) | C_2(k))$ for all $k \in M_{12}(\omega)$.

If $\exists k$ such that $E(U(f(C_1(k)), \cdot) | C_2(k)) < E(U(f(C_2(k)), \cdot) | C_2(k))$, then

$$\sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_1(k)), \cdot) | C_2(k)) < \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) | C_2(k))$$

that is:

$$P(M_{12}(\omega)) E(U(f(M_{12}(\omega)), \cdot) | M_{12}(\omega)) < \sum_{C_2(k) \subseteq M_{12}(\omega)} P(C_2(k)) E(U(f(C_2(k)), \cdot) | C_2(k))$$

which is a contradiction.

Hence we have $E(U(f(C_1(k)), \cdot) | C_2(k)) = E(U(f(C_2(k)), \cdot) | C_2(k))$ for all $k \in M_{12}(\omega)$.

As it is true for all ω , we have $E(U(f(C_1(k)), \cdot) | C_2(k)) = E(U(f(C_2(k)), \cdot) | C_2(k))$ for all $k \in \Omega$. The same reasoning applies for 2, 3 and 3, 1. \square

From **Result 4** and the assumption that all actions bring different expected utilities, we have

$$f(C_1(\omega)) = f(C_2(\omega)) = f(C_3(\omega)) \quad \forall \omega \in \Omega$$

Proof: [Proposition 1]:

Consider $X, X' \subseteq \Omega$, $X \cap X' = \emptyset$, such that $f(X) = f(X') = D^* \subseteq \mathcal{D}$. Let d^* be some element of D^* . Clearly, we have $E[U(d^*, \cdot) | X \cup X'] = \max_{d \in \mathcal{D}} E[U(d, \cdot) | X \cup X']$, which implies that $d^* \in f(X \cup X')$. Therefore, $D^* \subseteq f(X \cup X')$.

Let $d^{**} \in f(X \cup X')$. We have $E[U(d^{**}, \cdot) | X \cup X'] = \max_{d \in \mathcal{D}} E[U(d, \cdot) | X \cup X'] = E[U(d^*, \cdot) | X \cup X']$. If $d^{**} \notin D^*$, then $E[U(d^{**}, \cdot) | X] < E[U(d^*, \cdot) | X]$ and $E[U(d^{**}, \cdot) | X'] < E[U(d^*, \cdot) | X']$, which implies that one would have $E[U(d^{**}, \cdot) | X \cup X'] < E[U(d^*, \cdot) | X \cup X']$, which is a contradiction. Therefore, $d^{**} \in D^*$ and $F(X \cup X') \subseteq D^*$. \square

Proof: [Proposition 2]

Consider $X, X' \subseteq \Omega$, $X \cap X' = \emptyset$, such that $f(X) = f(X') = d^* \in \mathcal{D}$. Clearly, we have $\max_{d \in \mathcal{D}} E[U(d, \cdot) | X \cup X'] = E[U(d^*, \cdot) | X \cup X']$. Therefore, $d^* \in \operatorname{argmax}_{d \in \mathcal{D}} E[U(d, \cdot) | X \cup X']$. If $f(X \cup X') \neq d^*$, then it must be the case that $f(X \cup X') \succ d^*$. Yet if $f(X \cup X') \succ d^*$, then it must also be the case that $E[U(f(X \cup X'), \cdot) | X] < E[U(d^*, \cdot) | X]$ and that $E[U(f(X \cup X'), \cdot) | X'] < E[U(d^*, \cdot) | X']$. This implies that $E[U(f(X \cup X'), \cdot) | X \cup X'] < E[U(d^*, \cdot) | X \cup X']$, which is a contradiction. Therefore, $f(X \cup X') = d^*$, and f is union consistent.

Chapter 5

Communication, consensus and order. Who wants to speak first?

5.1. Introduction

Alice¹ and Bob are sitting in front of each other, both wearing either a red hat or a white hat. Suppose that the two hats are red. The teacher tells the children that there is at least one red hat, and asks them whether they know the color of their hat. The two children observe that the other's hat is red, but cannot infer the color of their own hat. The only way for them to answer the teacher is to communicate with each other. Suppose that Alice tells Bob that she does not know the color of her hat. Bob understands that his own hat is red, for if it had been white, Alice would have known that her hat was red. Now Bob knows the color of his hat. But then if he tells Alice that he knows the color of his hat, Alice will not learn anything, for the message of Bob would have been the same regardless of the color of her hat. Therefore, Alice has no interest to be the first to say whether she knows the color of her hat. This story illustrates the following fact. From the

¹This chapter is a joint work with Nicolas Houy.

moment that people communicate in order to be better informed, who gets to talk when is important: the communication process is not commutative, for different orders of speech may lead to different outcomes.

It is well known since Geanakoplos and Polemarchakis [1982] that in a group of rational agents, a process of simultaneous and public communication of posterior probabilities of an event leads to equality of all individual posteriors. Cave [1983] and Bacharach [1985] extended this result to simultaneous and public communication of decisions, assuming that the decision rule followed by agents satisfies a property called union consistency. Yet in most economic situations where agents have to speak with each other, communication is not simultaneous, and may not be public. It is common sense that each individual speaks one after the other according to a given protocol. Parikh and Krasucki [1990] considered the case where agents of a group communicate the value of some function f with each other, according to a pairwise protocol upon which they have agreed beforehand. They investigated what conditions on the function f and on the protocol guarantee that agents eventually reach a consensus, *i.e.* that from some stage on all the communicated values will be the same. They show that if the protocol is *fair*, that is if every participant receives information directly or indirectly from every other participant, and if the function f is *convex*, that is for all pair of disjoint events X, X' , there exists $a \in]0, 1[$ such that $f(X \cup X') = af(X) + (1-a)f(X')$, then communication will eventually lead to a consensus about the value of f .

The starting point of this work is to notice that, in Parikh and Krasucki's setting, different protocols may lead to different outcomes, in terms of consensus values of f as well as information learned by the agents during the communication process. In particular, it may well be the case that some agent learns more information when communicating according

to some protocol α than with some protocol β . It may also be the case that the most informative protocols are not the same for all agents. Therefore, if agents communicate in order to learn information from each other, they may disagree about the protocol they should use for communicating.

The issue we address in this chapter is the following. We consider the same setting as Parikh and Krasucki's, and we make the additional assumption that agents communicate with each other so as to learn information. Implicitly, agents are decision maker who try to be better informed in the sense of Blackwell [1953] so as to improve their decisions. An agent is said to be better informed with a structure I than with a structure I' if the maximum expected payoff yielded by I is larger than that yielded by I' for any payoff function and any prior probability. In a one-player decision problem, it has been shown that an agent is better informed with a partition Π than with a partition Π' if and only if Π is a refinement of Π' , namely if each cell of Π is included in a cell of Π' . Therefore, we say that an agent prefers a protocol α to a protocol β if and only if, at the end of the day, he has a finer partition when communicating with α than with β . The non-commutativity of the order of speech, as well as the fact that agents communicate so as to be better informed, imply that they have preferences over the set of possible orders of speech. Depending on the state of the world, Alice and Bob may prefer to speak first or second, or may be indifferent. If neither Alice nor Bob wants to speak first, communication can not take place. However, can we conclude that they will not learn anything from each other? The fact that Alice does not want to speak first is informative for Bob. Bob knows that if Alice knew the color of her hat, she wouldn't mind speaking first and saying that she knows the color of her hat. In this paper, we investigate what inferences can be made by rational agents from the common knowledge that some of them disagree about the order of speech.

The following situations are both possible. First, it can be common knowledge in a group of agents that some of them prefer the same order of speech. Second, it can be common knowledge in a group of agents that some of them prefer different orders of speech. However, we show the surprising result that if it is the case, then the consensus value of f must be the same whatever the order of speech. For instance, if it is common knowledge among Alice and Bob that they both want to speak first, then what they will communicate at the end of the day will be the same, whether Alice or Bob speaks first.

The chapter is organized as follows. In Section 2 we describe the model and the basic result of Parikh and Krasucki [1990]. Section 3 defines preferences over protocols and develops the result. Section 4 attempts to interpret and discuss our result. All proofs are given in the section 5.

5.2. Preliminary notions

Let Ω be the finite set of states of the world, and 2^Ω the set of possible events. There are N agents, each agent i being endowed with a partition Π_i of Ω . When the state $\omega \in \Omega$ occurs, agent i just knows that the true state of the world belongs to $\Pi_i(\omega)$, which is the cell of i 's partition that contains ω . We say that a partition Π is *finer* than a partition Π' if and only if for all ω , $\Pi(\omega) \subseteq \Pi'(\omega)$ and there exists ω' such that $\Pi(\omega') \subset \Pi'(\omega')$. A partition Π' is *coarser* than a partition Π if and only if Π is finer than Π' . The partition Π_i represents the ability of agent i to distinguish between the states of the world. The coarser his partition is, the less precise his information is, in the sense that he distinguishes among fewer states of the world. As usual, we say that an agent i endowed with a partition Π_i knows the event E at state ω if and only if $\Pi_i(\omega) \subseteq E$. We define the meet of the

partitions $\Pi_1, \Pi_2, \dots, \Pi_N$ as the finest common coarsening of these partitions, that is the finest partition M such that for all $\omega \in \Omega$ and for all $i = 1, \dots, N$, $\Pi_i(\omega) \subseteq M(\omega)$.

Before communicating, agents have to agree on a communication protocol that will be applied throughout the debate, and which determines which agents are allowed to speak at each date. We recall the formal definition of a communication protocol.

Definition 1 *A protocol α is a pair of functions (s, r) from \mathbb{N} to $2^{\{1, \dots, N\}}$. If $s(t) = S$ and $r(t) = R$, then we interpret S and R as, respectively, the set of senders and the set of receivers of the communication which takes place at time t .*

We note Γ the set of protocols. Along the debate, agents communicate by sending messages, which we assume to be delivered instantaneously, that is at time t , messages are simultaneously sent by every $i \in s(t)$ and heard by every $j \in r(t)$. As in Parikh and Krasucki [1990], we assume that the message sent is the private value of some function f defined from the set of subsets of Ω to \mathbb{R} . The private value of f for an agent i at state ω is $f(\Pi_i(\omega))$. This assumption implies that communication is completely non-strategic. This is obviously a strong assumption, especially in a setting where agents communicate in order to learn information. One implicitly assumes that agents are constrained to use such a communication rule, and behave like automatons.

Finally, the set of states of the world Ω , the individual partitions $(\Pi_i)_i$, and the message rule f define an *information model* $I = \langle \Omega, (\Pi_i)_i, f \rangle$.

Two assumptions are made on the protocol and on the function f to guarantee that iterated communication of the value of f leads to a consensus about the value of f . As in Parikh and Krasucki [1990], we assume that the protocol is *fair*. We adapt² Parikh and

²This definition is from Koessler [2001].

Krasucki's definition in our setting, but the meaning remains the same: a protocol is *fair* if and only if every participant in this protocol communicates directly or indirectly with every other participant infinitely many times. This condition is necessary so that nobody is excluded from communication.

Assumption 1 (A1) *The protocol α is fair, that is for all pair of players (i, j) , $i \neq j$, there exists an infinite number of finite sequences t_1, \dots, t_K , with $t_k \in \mathbb{N}$ for all $k \in \{1, \dots, K\}$, such that $i \in s(t_1)$ and $j \in r(t_K)$.*

Assumption 2 (A2) *f is convex, that is for all subsets $E, E' \subseteq \Omega$ such that $E \cap E' = \emptyset$, there exists $\alpha \in]0, 1[$ such that $f(E \cup E') = \alpha f(E) + (1 - \alpha)f(E')$.*

Note that we will have $f(E_1 \cup E_2 \cup \dots \cup E_k) = \sum_{i=1}^k \alpha_i f(E_i)$, with $\alpha_i \in]0, 1[\forall i$ and $\sum_{i=1}^k \alpha_i = 1$ provided that the E_i are pairwise disjoint events.

We now describe how information is aggregated during the debate. At a given date t , the senders $s(t)$ selected by the protocol (s, r) send a message heard by the receivers $r(t)$. Then each individual infers the set of states of the world that are compatible with the messages possibly sent, and updates his partition accordingly. Given an information model $\langle \Omega, (\Pi_i)_i, f \rangle$ and a communication protocol α , we define by induction on t the set $\Pi_i^\alpha(\omega, t)$ of possible states for an agent i at time t , given that the state of the world is ω :

$$\Pi_i^\alpha(\omega, 0) = \Pi_i(\omega) \text{ and for all } t \geq 1,$$

$$\Pi_i^\alpha(\omega, t+1) = \Pi_i^\alpha(\omega, t) \cap \{\omega' \in \Omega \mid f(\Pi_j^\alpha(\omega', t)) = f(\Pi_j^\alpha(\omega, t)) \forall j \in s(t)\} \text{ if } i \in r(t),$$

$$\Pi_i^\alpha(\omega, t+1) = \Pi_i^\alpha(\omega, t) \text{ otherwise.}$$

The next result states that for all i , for all ω , $f(\Pi_i^\alpha(\omega, t))$ has a limiting value, and

that this value does not depend on i . Under assumptions $A1$ and $A2$, participants in the protocol converge to a consensus about the value of f .

Proposition 1 (Parikh and Krasucki (1990)) *Let $\langle \Omega, (\Pi_i)_i, f \rangle$ be an information model, and α a communication protocol. Under assumptions $A1$ and $A2$, there exists a date T such that for all ω , for all i, j , and all $t, t' \geq T$, $f(\Pi_i^\alpha(\omega, t)) = f(\Pi_j^\alpha(\omega, t'))$.*

In the sequel, we will denote $\Pi_i^\alpha(\omega)$ the limiting value of $\Pi_i^\alpha(\omega, t)$, and Π_i^α will be called i 's partition of information at consensus. $f(\Pi^\alpha(\omega))$ will denote the limiting value of $f(\Pi_i^\alpha(\omega, t))$, which does not depend on i , and will be called the consensus value of f at state ω , given that the protocol is α .

5.3. Who wants to speak first? An agreement theorem.

We know from Parikh and Krasucki [1990] that given any protocol α , under assumptions $A1$ and $A2$, iterated communication of the private value of f eventually leads to a consensus about the value of f . The next proposition states that this value may vary according to the protocol.

Proposition 2 *There exist an information model $\langle \Omega, (\Pi_i)_i, f \rangle$ with f convex and two fair protocols α, β for which there exists ω such that $f(\Pi^\alpha(\omega)) \neq f(\Pi^\beta(\omega))$.*

This result can be proved easily for some union consistent functions f . However, to the best of our knowledge, it was not proved for conditional probabilities. As the posterior probabilities of an event are particular union consistent function, it could have been possible that there exist no information model with posterior probabilities such that order matters.

We exhibit an example where it does.

Example 1 Let $\Omega = \{1, \dots, 13\}$ be the set of states of the world. Suppose that Alice and Bob have a uniform prior P on Ω . They communicate in turn the private value of the function $f(\cdot) = P(\{2, 3, 4, 8, 12\} \mid \cdot)$, which is convex, and are endowed with the following partitions of Ω^3 :

$$\Pi_A = \{1, 3, 7, 8\}_{1/2}\{2, 6, 11, 12\}_{1/2}\{4, 5, 10\}_{1/3}\{9\}_0\{13\}_0$$

$$\Pi_B = \{1, 3, 5\}_{1/3}\{2\}_1\{4, 7, 9, 10, 12, 13\}_{1/3}\{6, 8\}_{1/2}\{11\}_0$$

If Alice speaks first (protocol α), individual partitions at consensus are:

$$\Pi_A^\alpha = \{1, 3, 7, 8\}_{1/2}\{2\}_1\{11\}_0\{6, 12\}_{1/2}\{4, 10\}_{1/2}\{5\}_0\{9\}_0\{13\}_0$$

$$\Pi_B^\alpha = \{1, 3\}_{1/2}\{5\}_0\{2\}_1\{4, 10\}_{1/2}\{7, 12\}_{1/2}\{9, 13\}_0\{6, 8\}_{1/2}\{11\}_0$$

If Bob speaks first (protocol β), individual partitions at consensus are:

$$\Pi_A^\beta = \{1, 3, 7\}_{1/3}\{8\}_1\{2\}_1\{6\}_0\{11\}_0\{12\}_1\{4, 5, 10\}_{1/3}\{9\}_0\{13\}_0$$

$$\Pi_B^\beta = \{1, 3, 5\}_{1/3}\{2\}_1\{4, 7, 10, \}_{1/3}\{12\}_1\{9, 13\}_0\{6\}_0\{8\}_1\{11\}_0$$

At state 1, the consensus value of f is $f(\{1, 3, 7, 8\}) = f(\{1, 3\}) = 1/2$ if Alice speaks first, whereas it is $f(\{1, 3, 7\}) = f(\{1, 3, 5\}) = 1/3$ if Bob speaks first.

We assume that agents are decision-makers who communicate with each other in order to be better informed. As a consequence, they prefer protocols that lead them to be better informed at the end of the day. What does «better informed» mean? Suppose that a decision maker has to choose an action, and that the set of payoff-relevant states of the

³The subscript reflects the posterior belief in each cell.

world is Θ . Let P and P' be two partitions of Θ . The decision maker is better informed with P than with P' if the maximal expected payoff yielded by P is larger than the maximal expected payoff yielded by P' , for any utility function, and any prior probability over Θ . It has been shown that the partition P is more informative than the partition P' if and only if P is a refinement of P' , namely if $P(\theta) \subseteq P'(\theta)$ for all $\theta \in \Theta$.

Before communication takes place, the set of states which are of matter of interest for agent i at state ω is $\Pi_i(\omega)$. Therefore, at state ω , agent i knows that he will be better informed with consensus partition Π_i^α than with consensus partition $\Pi_i^{\alpha'}$ if and only if Π_i^α induces a finer partition of $\Pi_i(\omega)$ than $\Pi_i^{\alpha'}$.

Definition 1 *Let α and α' be two protocols. An agent i is better informed with Π_i^α than with $\Pi_i^{\alpha'}$ at state ω if and only if*

- $\Pi_i^\alpha(\omega') \subseteq \Pi_i^{\alpha'}(\omega')$ for all $\omega' \in \Pi_i(\omega)$
- $\Pi_i^\alpha(\omega'') \subsetneq \Pi_i^{\alpha'}(\omega'')$ for some $\omega'' \in \Pi_i(\omega)$.

As a consequence, each agent knows *ex interim* which protocol he prefers among any two protocols if he is not indifferent.

Definition 2 (Preferences) *Let $I := \langle \Omega, (\Pi_i)_i, f \rangle$ be an information model, and α, β two protocols. The set of states of the world where agent i prefers α to β is denoted $B_i^I(\alpha, \beta)$ and is defined by*

$$B_i^I(\alpha, \beta) = \{\omega \in \Omega \mid \forall \omega' \in \Pi_i(\omega), \Pi_i^\alpha(\omega') \subseteq \Pi_i^\beta(\omega') \text{ and } \exists \omega'' \in \Pi_i(\omega) \text{ s.t. } \Pi_i^\alpha(\omega'') \subset \Pi_i^\beta(\omega'')\}$$

In Example 1, Alice and Bob are both better informed with the protocol α at state 4 and better informed with the protocol β at state 8. Hence at states 4 and 8, they agree on the protocol they prefer. On the contrary, at state 1, Alice and Bob end up strictly

better informed when they speak in second. What happens in that case? Suppose that state 1 occurs, and that Alice and Bob stand in front of each other waiting for the other to speak. Alice knows that the state of the world belongs to $\{1, 3, 7, 8\}$. She understands that the state of the world can not be 7 nor 8, for Bob would have spoken first at state 8 and would have been indifferent at state 7. Bob knows that the state of the world belongs to $\{1, 3, 5\}$. He understands that the state of the world can not be 5, for he knows that Alice prefers to speak first at state 5. Hence knowing that the other does not want to speak first makes Alice and Bob understand that the state of the world is in $\{1, 3\}$. From now, they have the same private information at state 1. As they cannot learn information from the communication process, they become indifferent between speaking first or second. This example addresses the question of whether it can be common knowledge among two persons that they disagree about the order of speech. More generally, what inferences can be made by rational agents of a group from the common knowledge that some of them disagree about the order of speech? Our main result states that if it is the case, then the consensus message is the same according to any protocol.

Theorem 1 *Let $I = \langle \Omega, (\Pi_i)_i, f \rangle$ be an information model such that A1 and A2 are satisfied, and α, β two protocols such that $\alpha \neq \beta$. Consider $a_1, a_2, b_1, b_2 \in \{\alpha, \beta\}$, with $a_1 \neq a_2$ and $b_1 \neq b_2$, and let us fix $i \neq j$. Assertions (1), (2) and (3) cannot be true simultaneously.*

(1) $B_i^I(a_1, a_2)$ and $B_j^I(b_1, b_2)$ are common knowledge at ω .

(2) $\omega \in B_i^I(a_1, a_2) \cap B_j^I(b_1, b_2)$ and $a_1 = b_2$.

(3) $f(\Pi^\alpha(\omega)) \neq f(\Pi^\beta(\omega))$.

The meaning of this result in the example described in introduction is the following.

- If (1) and (2) are true, namely if it is common knowledge at some state ω that Alice and Bob prefer to speak first, then (3) is false, *i.e.* the consensus value of f at ω is the same regardless of the person who speaks first.

- If (1) and (3) are true, namely if it is common knowledge at ω that Alice prefers $a_1 \in \{\alpha, \beta\}$ and Bob prefers $b_1 \in \{\alpha, \beta\}$, and if the consensus value of f differs according to whether the protocol is α or β , then (2) is false, *i.e.* Alice and Bob prefer the same protocol ($a_1 = b_1$).

- If (2) and (3) are true, namely if Alice and Bob prefer different orders of speech at ω , then (1) is false, *i.e.* these preferences are not common knowledge among them at ω .

The result of Theorem 1 is not due to the fact that assertions (1) and (2) or (1) and (3) or (2) and (3) are never true simultaneously.

Proposition 3 (i) *Assertions (1) and (2) of Theorem 2 can be true simultaneously.*

(ii) *Assertions (1) and (3) of Theorem 2 can be true simultaneously.*

(iii) *Assertions (2) and (3) of Theorem 2 can be true simultaneously.*

This proposition states that (i) it can be common knowledge among them that Alice and Bob prefer different orders of speech, (ii) it can be common knowledge among them that Alice and Bob prefer the same order of speech, and (iii) it is possible that Alice and Bob prefer different orders of speech which lead to different consensus values of f .

We prove point (i) with the following example, which describes a situation where it is common knowledge between Alice and Bob that both of them prefer to speak in second. The fact that both prefer to speak in second in order to be better informed is quite intuitive. When an agent is the second to speak, the first message he hears contains purely private

information of the other. When he is the first to speak, the first message he will hear will be a join of the other's private information and his private information, so he may not learn anything. However, we found another example which shows that there exist situations where both agents prefer to speak first. This example involves 288 states of the world, so we did not include it in the thesis.

Example 2 *The set of states of the world is $\Omega = \{1, 2, 3, 4, 5, 6, 7\}$ and Alice and Bob are endowed with a uniform prior P on Ω . They communicate in turn the private value of the function $f(\cdot) = P(\{1, 2, 7\} \mid \cdot)$ and are endowed with the following partitions:*

$$\Pi_A = \{1, 2\}_1\{3, 4\}_0\{5, 6, 7\}_{1/3}$$

$$\Pi_B = \{1, 7\}_1\{2, 3, 6\}_{1/3}\{4, 5\}_0$$

If Alice speaks first (protocol α), individual partitions at consensus are:

$$\left\{ \begin{array}{l} \Pi_A^\alpha = \{1, 2\}_1\{3, 4\}_0\{5, 6\}_0\{7\}_1 \\ \Pi_B^\alpha = \{1\}_1\{2\}_1\{3\}_0\{4\}_0\{5\}_0\{6\}_0\{7\}_1 \end{array} \right.$$

If Bob speaks first (protocol β), individual partitions at consensus are:

$$\left\{ \begin{array}{l} \Pi_A^\beta = \{1\}_1\{2\}_1\{3\}_0\{4\}_0\{5\}_0\{6\}_0\{7\}_1 \\ \Pi_B^\beta = \{1, 7\}_1\{2\}_1\{3, 6\}_0\{4, 5\}_0 \end{array} \right.$$

At every state of the world, Alice and Bob both prefer to speak in second: $B_A(\beta, \alpha) = \Omega$ and $B_B(\alpha, \beta) = \Omega$, hence at every state of the world, it is common knowledge among Alice and Bob that Alice prefers the order β and Bob the order α . However, it does not contradict Theorem 2 as for all ω , $f(\Pi^\alpha(\omega)) = f(\Pi^\beta(\omega))$.

We prove point (ii) with the following example, which shows that it is possible that both agents prefer the same order of speech.

Example 3 *The set of states of the world is $\Omega = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$ and Alice and Bob are endowed with a uniform prior P on Ω . They communicate in turn the private value of the function $f(\cdot) = P(\{1, 6, 7, 9\} \mid \cdot)$ and are endowed with the following partitions:*

$$\Pi_A = \{1, 2, 4, 5, 9\}_{2/5} \{3, 6, 7, 8\}_{1/2}$$

$$\Pi_B = \{1, 3, 7\}_{1/3} \{2, 5, 8\}_0 \{4, 6, 9\}_{2/3}$$

If Alice speaks first (protocol α), individual partitions at consensus are:

$$\left\{ \begin{array}{l} \Pi_A^\alpha = \{1\}_1 \{2, 5\}_0 \{4, 9\}_{1/2} \{3, 7\}_{1/2} \{6\}_1 \{8\}_0 \\ \Pi_B^\alpha = \{1\}_1 \{2, 5\}_0 \{4, 9\}_{1/2} \{3, 7\}_{1/2} \{6\}_1 \{8\}_0 \end{array} \right.$$

If Bob speaks first (protocol β), individual partitions at consensus are:

$$\left\{ \begin{array}{l} \Pi_A^\beta = \{1, 4, 9\}_{2/3} \{2, 5\}_0 \{3, 6, 7\}_{2/3} \{8\}_0 \\ \Pi_B^\beta = \{1, 3, 7\}_{2/3} \{2, 5, 8\}_0 \{4, 6, 9\}_{2/3} \end{array} \right.$$

At every state of the world, Alice and Bob prefer that Alice speaks first: $B_A(\alpha, \beta) = B_B(\alpha, \beta) = \Omega$, hence it is common knowledge at any state that both prefer the order α .

Finally, we prove point (iii) with Example 1 in section 2. The partition of common knowledge is $M = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13\}$. At state 1, Alice and Bob prefer to speak second, and $f(\Pi^\alpha(1)) = 1/3 \neq f(\Pi^\beta(1)) = 1/2$. However, this is not *common knowledge*, for Bob prefers to speak first at states 6 and 8.

5.4. Concluding remarks

What should be retained from Theorem 1 is that it cannot be common knowledge in a group of agents that two of them disagree about the protocol they should use to

communicate about f , if the consensus value of f actually differs according to each protocol. This result is somewhat surprising, as it states basically that common knowledge about fineness of partitions has implications in terms of messages. Indeed, the result is not due to the fact that common knowledge that two agents disagree about two protocols α and β implies that the partition of common knowledge is the same with α and β . Formally, we have:

Proposition 4 *Let $I = \langle \Omega, (\Pi_i)_i, f \rangle$ be an information model such that A1 and A2 are satisfied, and α, β two protocols such that $\alpha \neq \beta$. Consider $a_1, a_2, b_1, b_2 \in \{\alpha, \beta\}$, with $a_1 \neq a_2$ and $b_1 \neq b_2$, and let us fix $i \neq j$. Assertions (1), (2) and (3') can be true simultaneously.*

(1) $B_i(a_1, a_2)$ and $B_j(b_1, b_2)$ are common knowledge at ω

(2) $\omega \in B_i(a_1, a_2) \cap B_j(b_1, b_2)$ and $a_1 = b_2$

(3') $\Pi^\alpha(\omega) \neq \Pi^\beta(\omega)$

We prove it with **Example 2**, in which it is common knowledge in every state of the world that Alice and Bob both want to speak second (which implies that assertions (1) and (2) are true). The partition of common knowledge at consensus is $\Pi^\alpha : \{1, 2\}\{3, 4\}\{5, 6\}\{7\}$ if Alice speaks first, and is $\Pi^\beta : \{1, 7\}\{2\}\{3, 6\}\{4, 5\}$ if Bob speaks first. However, we have $f(\Pi^\alpha(\omega)) = f(\Pi^\beta(\omega))$ for all ω even if $\Pi^\alpha(\omega) \neq \Pi^\beta(\omega)$ for all ω .

A first limit to this result is that conditions of application of Theorem 1 are quite strong, in particular Assertion 2. Indeed, because of the way we defined them, preferences over protocols are not complete. Therefore, the likelihood of a situation where two agents would disagree between two protocols could be small.

Another obvious limit of our result is that the justification we used for such preferences over protocols works only in one-player decision problems. Indeed, more information does not always mean better information in interactive situations. We know by Scarsini and Zamir [1997], or Kamien, Tauman and Zamir [1990] among others, that in some games, players might prefer dropping some payoff-relevant information because their equilibrium payoff would then be higher. However, in particular games such as zero-sum games and common interest games, each player improves his expected payoff by learning more information.

5.5. Proof

Consider an information model $I = \langle \Omega, (\Pi_i)_i, f \rangle$, and α, β two protocols such that $\alpha \neq \beta$. Let us show that if points 1) and 2) of theorem 1 are true, then point 3) is false. We show that if there exist two agents i, j and a state ω such that $B_i^I(\alpha, \beta)$ and $B_j^I(\beta, \alpha)$ are common knowledge at ω , then $f(\Pi^\alpha(\omega)) = f(\Pi^\beta(\omega))$. Clearly, the proof still holds if we invert α and β .

Recall that $M(\omega)$ denotes the meet of individual partitions before communication takes place: $M = \bigwedge_{i=1}^n \Pi_i$. We note Π^α the meet of the individual partitions at consensus, given that the protocol is α : $\Pi^\alpha = \bigwedge_{i=1}^n \Pi_i^\alpha$.

If $B_i^I(\alpha, \beta)$ and $B_j^I(\beta, \alpha)$ are common knowledge at ω , then we have

$$M(\omega) \subseteq B_i(\alpha, \beta) \cap B_j(\beta, \alpha)$$

As $\Pi^\alpha(\omega) \subseteq M(\omega)$ and $\Pi^\beta(\omega) \subseteq M(\omega) \forall \omega$, we have $\Pi^\alpha(\omega) \cap \Pi^\beta(\omega) \subseteq M(\omega) \forall \omega$.

Hence we have

$$\Pi^\alpha(\omega) \cap \Pi^\beta(\omega) \subseteq B_i^I(\alpha, \beta) \cap B_j^I(\beta, \alpha) \tag{5.1}$$

Consider some $\omega' \in \Pi^\alpha(\omega) \cap \Pi^\beta(\omega)$ (which is not empty as $\omega \in \Pi^\alpha(\omega) \cap \Pi^\beta(\omega)$). By definition of the meet, we have $\Pi_i^\alpha(\omega') \subseteq \Pi^\alpha(\omega')$ and $\Pi_i^\beta(\omega') \subseteq \Pi^\beta(\omega')$. As $\omega' \in \Pi^\alpha(\omega) \cap \Pi^\beta(\omega)$, we have $\Pi^\alpha(\omega') = \Pi^\alpha(\omega)$ and $\Pi^\beta(\omega') = \Pi^\beta(\omega)$. Then we have

$$\Pi_i^\alpha(\omega') \subseteq \Pi^\alpha(\omega) \text{ and } \Pi_i^\beta(\omega') \subseteq \Pi^\beta(\omega) \quad (5.2)$$

By (5.1), $\omega' \in B_i^I(\alpha, \beta)$. It implies that $\Pi_i^\alpha(\omega') \subseteq \Pi_i^\beta(\omega')$. Yet $\Pi_i^\beta(\omega') \subseteq \Pi^\beta(\omega)$ by (5.2). Then we have

$$\Pi_i^\alpha(\omega') \subseteq \Pi^\alpha(\omega) \cap \Pi^\beta(\omega)$$

As this is true for every $\omega' \in \Pi^\alpha(\omega) \cap \Pi^\beta(\omega)$, we have

$$\Pi^\alpha(\omega) \cap \Pi^\beta(\omega) = \bigcup_{\omega' \in \Pi^\alpha(\omega) \cap \Pi^\beta(\omega)} \Pi_i^\alpha(\omega')$$

By Proposition 1 of Parikh and Krasucki [1990], $\forall i, j, f(\Pi_i^\alpha(\omega)) = f(\Pi_j^\alpha(\omega))$ for all ω . By definition of the meet, it implies that $\forall \omega' \in \Pi^\alpha(\omega), f(\Pi_i^\alpha(\omega')) = f(\Pi_i^\alpha(\omega))$. As f is convex, it is also union consistent, then we have $f(\Pi^\alpha(\omega) \cap \Pi^\beta(\omega)) = f(\Pi^\alpha(\omega))$.

The same reasoning applied to $\Pi_j^\beta(\omega)$ boils down to $f(\Pi^\alpha(\omega) \cap \Pi^\beta(\omega)) = f(\Pi^\beta(\omega))$.

Hence $f(\Pi^\alpha(\omega)) = f(\Pi^\beta(\omega)) \square$

Chapter 6

Comparative study of some properties of decision rules

6.1. Introduction

Aumann's result about the impossibility of agreeing to disagree has given rise to a vast literature, dealing both with the implications of common knowledge of events for economic behavior, and with the emergence of common knowledge situations in communication protocols. In any case, agents follow decision rules which prescribe what action to make as a function of any information situation they might be in. The impossibility of agreeing to disagree requires that individual decision rules satisfy some conditions. Cave [1983] and Bacharach [1985] need union consistency to show that common knowledge of individual decisions negates asymmetric information; Parikh and Krasucki [1990] need convexity to show that pairwise communication of individual decisions eventually leads to consensus in decisions; Aumann [1976], Geanakoplos and Polemarchakis [1982], McKelvey and Page [1986] use posterior probabilities, Sebenius and Geanakoplos [1983] and Nielsen *et al.* [1990] use conditional expectations, which are both convex functions. In Chapter 3, we use balanced

union consistency to show that common knowledge of an exhaustive statistic of individual decisions implies equality of decisions. In Chapter 4, we show that communication with an argmax rule eventually leads to consensus in any fair protocol.

The relation between convexity, weak convexity and union consistency has been made clear by Parikh and Krasucki [1990]: the set of convex functions is strictly included in the set of weakly convex functions, which is also strictly included in the set of union consistent functions. In this chapter, we first want to examine how the new conditions we introduced in the literature, balanced union consistency and argmax, relate to convexity, weak convexity and union consistency. Second, we discuss whether each condition is sufficient to guarantee that 1) common knowledge of an exhaustive statistic of individual decisions implies equality of decisions, and that 2) communication according to a fair protocol eventually leads to consensus. We show in Chapter 3 that common knowledge of an exhaustive statistic may not imply consensus with union consistent decision rules. What about argmax, weakly and simply convex rules? Parikh and Krasucki [1990] show that consensus may fail to occur with union consistency and weak convexity. What about balanced union consistency?

6.2. Relation between union consistency, weak convexity, convexity, balanced union consistency and argmax

Let us first recall the definition of each condition.

- **Union consistency:** f is union consistent if $\forall E, F \subseteq \Omega$, such that $E \cap F = \emptyset$
 $f(E) = f(F) \Rightarrow f(E \cup F) = f(E) = f(F)$.

We denote **UC** the set of union consistent functions.

- **Balanced union consistency:** f is balanced union consistent if for all balanced family¹ \mathcal{B} of Ω , $f(S) = d \forall S \in \mathcal{B} \Rightarrow f(\bigcup_{S \in \mathcal{B}} S) = d$.

We denote \mathbf{B} the set of balanced union consistent functions.

- **Argmax:** f is an argmax function if there exist a function $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$ and a probability P over Ω such that $\forall X \subseteq \Omega$, $f(X) = \operatorname{argmax}_{d \in \mathcal{D}} E[U(d, \cdot) | X]$.

We denote \mathbf{A} the set of argmax functions.

- **Weak convexity:** f is weakly convex if $\forall E, F \subseteq \Omega$ such that $E \cap F = \emptyset$, $\exists \alpha \in [0, 1]$ such that $f(E \cup F) = \alpha f(E) + (1 - \alpha)f(F)$.

We denote \mathbf{WC} the set of weakly convex functions.

- **Convexity:** f is convex if $\forall E, F \subseteq \Omega$ such that $E \cap F = \emptyset$, $\exists \alpha \in]0, 1[$ such that $f(E \cup F) = \alpha f(E) + (1 - \alpha)f(F)$.

We denote \mathbf{C} the set of convex functions.

Balanced union consistent functions, argmax function, weakly and strictly convex functions are all union consistent. Parikh and Krasucki [1990] show that convexity implies weak convexity which implies union consistency ($\mathbf{C} \subseteq \mathbf{WC} \subseteq \mathbf{U}$). We show in chapter 3 that argmax functions are balanced union consistent, and that balanced union consistency implies union consistency ($\mathbf{A} \subseteq \mathbf{B} \subseteq \mathbf{U}$). However, there is no inclusion relation between the sets of weakly or simply convex functions and balanced union consistent functions, as well as there is no inclusion relation between the sets of weakly or simply convex functions and argmax functions.

¹A non-empty family $\mathcal{B} \subseteq 2^\Omega$ is balanced if there exists a family of non-negative reals $\{\lambda_S\}_{S \in \mathcal{B}}$, called balancing coefficients, such that $\sum_{S \in \mathcal{B}, \omega \in S} \lambda_S = 1$ for every $\omega \in \bigcup_{S \in \mathcal{B}} S$.

Proposition We have:

1. $\mathbf{C} \subsetneq \mathbf{WC} \subsetneq \mathbf{UC}$,
2. $\mathbf{A} \subsetneq \mathbf{B} \subsetneq \mathbf{UC}$,
3. $\mathbf{A} \cap \mathbf{C} \neq \emptyset$, but $\mathbf{A} \not\subseteq \mathbf{WC}$ and $\mathbf{C} \not\subseteq \mathbf{A}$,
4. $\mathbf{B} \cap \mathbf{C} \neq \emptyset$, but $\mathbf{B} \not\subseteq \mathbf{WC}$ and $\mathbf{C} \not\subseteq \mathbf{B}$

We show points 3 and 4 with three examples. The next example shows that there exists a convex and balanced union consistent function which cannot be represented as the argmax of an expected utility.

Example 1 ($\mathbf{C} \not\subseteq \mathbf{A}$, $\mathbf{B} \not\subseteq \mathbf{A}$) *Let the set of states of the world be $\Omega = \{1, 2, 3, 4, 5\}$, and the function f defined from 2^Ω into $[0, 1]$ by:*

$$f(\{1\}) = f(\{2\}) = f(\{1, 2\}) = 1, \quad f(\{3\}) = f(\{4\}) = f(\{5\}) = f(\{3, 4\}) = f(\{3, 5\}) = f(\{4, 5\}) = f(\{3, 4, 5\}) = 0$$

$$f(\{1, 3\}) = f(\{1, 4\}) = f(\{1, 5\}) = f(\{2, 3\}) = f(\{2, 4\}) = f(\{2, 5\}) = f(\{1, 2, 3, 4\}) = f(\{1, 2, 3, 5\}) = f(\{1, 2, 4, 5\}) = 1/2$$

$$f(\{1, 4, 5\}) = f(\{2, 3, 4\}) = 1/3^*$$

$$f(\{1, 3, 4\}) = f(\{2, 4, 5\}) = 1/4^{**}$$

$$f(\{1, 3, 5\}) = f(\{2, 3, 5\}) = 1/4$$

$$f(\{1, 2, 3\}) = f(\{1, 2, 4\}) = f(\{1, 2, 5\}) = 3/4$$

$$f(\{1, 3, 4, 5\}) = f(\{2, 3, 4, 5\}) = 1/6$$

$$f(\{1, 2, 3, 4, 5\}) = 2/5$$

Such a function f is strictly convex and balanced union consistent. However, there exist no utility function $U : [0, 1] \times \Omega \rightarrow \mathbb{R}$, and no probability distribution P such that f is the

argmax of the conditional expectation of U . Suppose that there exist a utility function U and a probability P such that $f(X) = \operatorname{argmax}_{d \in [0,1]} E[U(d, \cdot) | X]$ for all X . Then by * we have:

$$P(1)U(1/3, 1) + P(4)U(1/3, 4) + P(5)U(1/3, 5) > P(1)U(1/4, 1) + P(4)U(1/4, 4) + P(5)U(1/4, 5)$$

and by ** we have:

$$P(1)U(1/4, 1) + P(3)U(1/4, 3) + P(4)U(1/4, 4) > P(1)U(1/3, 1) + P(3)U(1/3, 3) + P(4)U(1/3, 4)$$

which implies that

$$P(3)U(1/4, 3) + P(5)U(1/3, 5) > P(3)U(1/3, 3) + P(5)U(1/4, 5) \quad (6.1)$$

Moreover, by * we have:

$$P(2)U(1/3, 2) + P(3)U(1/3, 3) + P(4)U(1/3, 4) > P(2)U(1/4, 2) + P(3)U(1/4, 3) + P(4)U(1/4, 4)$$

and by ** we have:

$$P(2)U(1/4, 2) + P(4)U(1/4, 4) + P(5)U(1/4, 5) > P(2)U(1/3, 2) + P(4)U(1/3, 4) + P(5)U(1/3, 5)$$

which implies that

$$P(3)U(1/4, 3) + P(5)U(1/3, 5) < P(3)U(1/3, 3) + P(5)U(1/4, 5) \quad (6.2)$$

(6.1) and (6.2) together bring the contradiction.

We now show that there exists an argmax function (and therefore, a balanced union consistent function), which is not weakly convex.

Example 2 (A $\not\subseteq$ WC, B $\not\subseteq$ WC) Consider the case where $\Omega = \{1, 2, 3, 4\}$, $\mathcal{D} = \{a, b, c\}$,

P is uniform ($P(\omega) = 1/4 \forall \omega$) and the utility function $U : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$ is defined by:

$$U(a, 1) = 1, U(a, 2) = 0, U(a, 3) = 1, U(a, 4) = 0$$

$$U(b, 1) = 0, U(b, 2) = 1, U(b, 3) = 2/3, U(b, 4) = 2/3$$

$$U(c, 1) = 2/3, U(c, 2) = 2/3, U(c, 3) = 0, U(c, 4) = 1$$

Consider the argmax function $f : 2^\Omega \rightarrow \mathcal{D}$ defined by $f(X) = \operatorname{argmax}_{d \in \mathcal{D}} E[U(d, \cdot) \mid X]$.

We have in particular:

$$f(\{1\}) = a, f(\{2\}) = b, f(\{3\}) = a, f(\{4\}) = c, f(\{1, 2\}) = c, f(\{3, 4\}) = b$$

Let us show that f is not weakly convex, namely that there exist no function $g : \mathcal{D} \rightarrow \mathbb{R}$, such that $g \circ f : 2^\Omega \rightarrow \mathbb{R}$ is weakly convex.

For any one to one function $g : \mathcal{D} \rightarrow \mathbb{R}$, six cases are possible. We show that in each case, $g \circ f$ is not weakly convex.

1. If $g(a) < g(b) < g(c)$, then $g \circ f(\{1\}) < g \circ f(\{2\}) < g \circ f(\{1, 2\})$.
2. If $g(a) < g(c) < g(b)$, then $g \circ f(\{3\}) < g \circ f(\{4\}) < g \circ f(\{3, 4\})$.
3. If $g(b) < g(a) < g(c)$, then $g \circ f(\{3, 4\}) < g \circ f(\{3\}) < g \circ f(\{4\})$.
4. If $g(b) < g(c) < g(a)$, then $g \circ f(\{3, 4\}) < g \circ f(\{4\}) < g \circ f(\{3\})$.
5. If $g(c) < g(a) < g(b)$, then $g \circ f(\{1, 2\}) < g \circ f(\{1\}) < g \circ f(\{2\})$.
6. If $g(c) < g(b) < g(a)$, then $g \circ f(\{1, 2\}) < g \circ f(\{2\}) < g \circ f(\{1\})$.

This example proves that argmax functions (and then, balanced union consistent functions), may not be weakly convex.

Finally, we show that convexity does not imply balanced union consistency. The example of a convex function which is not balanced union consistent is far from trivial. We

conjecture that such an example requires an even number of states of the world, and four states of the world are not enough.

Example 3 (C $\not\subseteq$ B) Let the set of states of the world be $\Omega = \{1, 2, 3, 4, 5, 6\}$, and $f :$

$2^\Omega \setminus \emptyset \rightarrow \mathbb{R}$ be defined by:

$$f(\{1\}) = f(\{3\}) = f(\{1, 3\}) = 1$$

$$f(\{1, 3, 4\}) = f(\{1, 3, 6\}) = 2$$

$$f(\{3, 4\}) = f(\{1, 4\}) = 3$$

$$f(\{1, 3, 4, 6\}) = 4$$

$$f(\{1, 6\}) = f(\{3, 6\}) = 5$$

$$f(\{3, 4, 6\}) = f(\{1, 4, 6\}) = 6$$

$$f(\{4\}) = f(\{6\}) = f(\{4, 6\}) = 7$$

$$f(\{1, 2, 3, 4, 6\}) = f(\{1, 3, 4, 5, 6\}) = 8$$

$$f(\{1, 2, 3, 6\}) = f(\{1, 4, 5, 6\}) = f(\{1, 3, 5, 6\}) = f(\{1, 2, 4, 6\}) = f(\{3, 4, 5, 6\}) = f(\{1, 3, 4, 5\}) =$$

$$f(\{2, 3, 4, 6\}) = f(\{1, 2, 3, 4\}) = 9$$

$$f(\{1, 2, 3\})^* = f(\{1, 4, 5\})^* = f(\{2, 4, 6\})^* = f(\{3, 5, 6\})^* = f(\{1, 2, 6\}) = f(\{1, 5, 6\}) = 10$$

$$f(\{1, 2, 3, 4, 5, 6\}) = 11$$

$$f(\{1, 2, 4\}) = f(\{1, 3, 5\}) = f(\{2, 3, 6\}) = f(\{4, 5, 6\}) = f(\{3, 4, 5\}) = f(\{2, 3, 4\}) = 12$$

$$f(\{1, 2, 3, 4, 5\}) = f(\{1, 2, 3, 5, 6\}) = f(\{1, 2, 4, 5, 6\}) = f(\{2, 3, 4, 5, 6\}) = 13$$

$$f(\{2, 3\}) = f(\{2, 4\}) = f(\{3, 5\}) = f(\{4, 5\}) = f(\{2, 3, 4, 5\}) = 14$$

$$f(\{1, 2, 3, 5\}) = f(\{1, 2, 4, 5\}) = f(\{2, 3, 5, 6\}) = f(\{2, 4, 5, 6\}) = 15$$

$$f(\{1, 2\}) = f(\{1, 5\}) = f(\{2, 6\}) = f(\{5, 6\}) = f(\{1, 2, 5, 6\}) = 16$$

$$f(\{2, 3, 5\}) = f(\{2, 4, 5\}) = f(\{1, 2, 5\}) = f(\{2, 5, 6\}) = 17$$

$$f(\{2\}) = f(\{5\}) = f(\{2, 5\}) = 18$$

Such a function f is strictly convex, but is not balanced union consistent. Indeed, the family $B = \{\{1, 2, 3\}, \{1, 4, 5\}, \{2, 4, 6\}, \{3, 5, 6\}\}$ is balanced with respect to coefficients $1/2$ for each element of B . If f were balanced union consistent, the fact that $f(\{1, 2, 3\}) = f(\{1, 4, 5\}) = f(\{2, 4, 6\}) = f(\{3, 5, 6\}) = 10$ would imply that $f(\{1, 2, 3, 4, 5, 6\}) = 10$, whereas $f(\{1, 2, 3, 4, 5, 6\}) = 11$.

6.3. Relation between the results

We showed in chapter 3 that common knowledge of an exhaustive statistic of individual decisions implies consensus if the decision rule is balanced union consistent. We also show with a counter example that union consistency is not sufficient to guarantee the consensus result, if the number of agents is greater than 4. As argmax functions are particular balanced union consistent functions, the result of chapter 3 still holds for argmax decision rules. However, some convex functions are not balanced union consistent. We show with the next example that for those functions, common knowledge of an exhaustive statistic of individual decisions may not imply consensus.

Example 4 (No consensus with an exhaustive statistic in C) *Let the set of states of the world be $\Omega = \{1, 2, 3, 4, 5, 6\}$ and consider the convex function f defined as in Example 3l. Consider four agents, endowed with the following partitions:*

$$\Pi_1 : \{1, 2, 3\}_{10}\{4, 5, 6\}_{12}$$

$$\Pi_2 : \{1, 4, 5\}_{10}\{2, 3, 6\}_{12}$$

$$\Pi_3 : \{1, 3, 5\}_{12}\{2, 4, 6\}_{10}$$

$$\Pi_4 : \{1, 2, 4\}_{12}\{3, 5, 6\}_{10}$$

At every state of the world, two out of four agents take the decision 10, and two out of

four the decision 12. Therefore, the statistic $\Phi^*(\omega)$ is common knowledge at every state ω , without implying consensus.

Parikh and Krasucki [1990] show that communication in a fair protocol eventually leads to consensus if the function whose values are communicated is convex. They show that the result does not hold for weakly convex functions if the number of agents is larger or equal to 4, and does not hold neither for union consistent functions for more than two agents. We showed in chapter 4 that consensus emerge in any fair protocol if the functions whose values are communicated is the argmax of an expected utility. Argmax functions are balanced union consistent, but some balanced union consistent functions are not the argmax of an expected utility. We show with the following example, given in Parikh and Krasucki [1990, p 184], that consensus may fail to occur in some fair protocols for non-argmax, but balanced union consistent functions.

Example 5 (No consensus emerge in B with a fair protocol) *The set of states of the world is $\Omega = \{1, 2, 3, 4, 5, 6, 7, 8\}$. Consider the function $f : 2^\Omega \setminus \emptyset \rightarrow \mathbb{N}$ defined as follows: let the elements of $2^\Omega \setminus \emptyset$ be numbered X_1, X_2, \dots , and let $\text{num}(X_i)$ be i . Now let*

$$f(\{1, 2\}) = f(\{3, 4\}) = f(\{1, 2, 3, 4\}) = 1$$

$$f(\{5, 6\}) = f(\{7, 8\}) = f(\{5, 6, 7, 8\}) = 2$$

$$f(\{1, 3\}) = f(\{5, 7\}) = f(\{1, 3, 5, 7\}) = 3$$

$$f(\{2, 4\}) = f(\{6, 8\}) = f(\{2, 4, 6, 8\}) = 4$$

$$f(\{1, 5\}) = f(\{2, 6\}) = f(\{1, 2, 5, 6\}) = 5$$

$$f(\{3, 7\}) = f(\{4, 8\}) = f(\{3, 4, 7, 8\}) = 6$$

and for all other elements X of $2^\Omega \setminus \emptyset$, we let $f(X) = \text{num}(X) + 6$. This ensures that f has both the union consistency and the balanced union consistency properties.

Now consider three agents, endowed with the following partitions:²

$$\Pi_A : \{1, 2\}_1 \{3, 4\}_1 \{5, 6\}_2 \{7, 8\}_2$$

$$\Pi_B : \{1, 3\}_3 \{2, 4\}_4 \{5, 7\}_3 \{6, 8\}_4$$

$$\Pi_C : \{1, 5\}_5 \{2, 6\}_5 \{3, 7\}_6 \{4, 8\}_6$$

Agent A speaks to agent B , who speaks to agent C , who speaks to agent A , and so on.

The partition of common knowledge among agents A and B is

$$M^{AB} : \{1, 2, 3, 4\} \{5, 6, 7, 8\}$$

As agent A sends the message 1 at states 1, 2, 3 and 4, and the message 2 at states 5, 6, 7 and 8, A 's message is common knowledge among A and B at every state of the world. Therefore, agent B does not learn anything from A 's message. The partition of common knowledge among agents B and C is

$$M^{BC} : \{1, 3, 5, 7\} \{2, 4, 6, 8\}$$

The set of states where B sends the message 3 is $\{1, 3, 5, 7\}$, and the set of states where B sends the message 4 is $\{2, 4, 6, 8\}$. Again, agent C does not learn anything from B 's message. Finally, the partition of common knowledge among agents C and A is

$$M^{AC} : \{1, 2, 5, 6\} \{3, 4, 7, 8\}$$


The set of states where C sends the message 5 is $\{1, 2, 5, 6\}$, and the set of states where C sends the message 6 is $\{3, 4, 7, 8\}$. As a consequence, agent A does not learn anything from C 's message.

²The subscript reflects the decision associated to each cell.

This example shows that consensus may fail to emerge in a fair protocol, if the function f is balanced union consistent (and then union consistent).

6.4. Conclusion

To conclude, we sum up the relations between union consistency, balanced union consistency, argmax, convexity and weak convexity in the following figure.



Among the five conditions we identified, convexity and argmax are the only one ensuring that communication of individual decisions leads to consensus in any fair protocol, for more than three agents. Balanced union consistency is the only one ensuring that common knowledge of an exhaustive statistic of individual decisions implies a consensus, for more than three agents.

Conclusion

This thesis is about common knowledge and some of its implications. Common knowledge of an event is the particular state of interactive knowledge where everybody knows the event, everybody knows that everybody knows the event, and so on *ad infinitum*. Common knowledge is inherent to much of social life, and is crucial for coordination problems, which are situations in which each person wants to participate to a joint action only if others participate also. Another implication of common knowledge has been illuminated by Aumann [1976] in his celebrated article Agreeing to Disagree. He showed that if two rational agents have the same prior probability, then they cannot have commonly known differences in posterior probabilities, despite different conditioning information. This result suggested that asymmetric information had less explanatory power than might be thought: in the absence of differences in prior beliefs, asymmetric information cannot explain commonly known differences in posterior beliefs. Aumann's theorem gave rise to a literature called the Agreeing to Disagree literature, which addresses basically the same question: to what extent differences in beliefs and decisions can be explained by asymmetric information? The most obvious contribution of this literature to economic theory are no-trade theorems. These theorems state that trade among rational agents cannot be explained on the basis of asymmetric information. To restore the conventional understanding of speculation and trade, one has to assume either that traders are boundedly rational, or that they hold

different prior probabilities.

This thesis is a contribution to the Agreeing to Disagree literature. It is made of six chapters. The first two chapters have introduced background material and have surveyed the literature. The three next chapters were original contributions. The last chapter compared various conditions of the three contributions to conditions that can be found in the literature.

In Chapter 1, we presented Aumann structures, which is the model of knowledge used in the Agreeing to Disagree literature and in this thesis. We insisted on two controversial issues. The first one deals with what agents know about the others' knowledge. We made an attempt to answer the questions of whether individual partitions are "common knowledge" to all agents, and whether some "common knowledge" of individual partitions was required for the meet to be the partition of common knowledge. Going back to epistemic logic, we emphasized that knowledge is part of the description of the states, that partitions are only representation tool, and that it makes no sense, to the best of our understanding, wondering whether such representation tool are "known" by agents. The second controversial issue deals with updating in Aumann structures, and is stated as follows. States of the world describe individual knowledge. When agents update their information, their knowledge change. Therefore, we may wonder whether states of the world change when agents update their information. Following some recent works of Bonanno [2004] and Board [2004], we emphasize that states do not change if they describe individual knowledge at each date. Therefore, the revision rule used in the literature is effectively a way of representing new knowledge of the same uncertainty.

In Chapter 2, we reviewed the Agreeing to Disagree literature. We classified the results in two groups. Results in the first group answer the following question. Under what

conditions common knowledge of a statistic of individual decisions implies that these decisions do not reflect the differential information that each agent possesses? Result in the second group answer the question of the emergence of common knowledge of individual decisions. Under what conditions communication of individual decisions leads to common knowledge of decisions? We particularly insisted on the difference between public and non-public communication protocols, for the way individual decisions become common knowledge essentially differs whether the protocol is public or not. In public protocols, common knowledge of individual decisions emerge without any restrictions on individual decision rules. In particular, commonness of decision rules is not required. In non-public protocols however, commonness of decision rules is required because common knowledge of individual decisions emerge *via* the consensus.

In Chapter 3, we provided an answer to the first question addressed in the literature, in the case where decisions may not be posterior probabilities of some event. We showed that if the statistic is exhaustive, and if individual decision rules are balanced union consistent, then common knowledge of a statistic of individual decisions negates asymmetric information. Exhaustiveness imposes that the statistic should describe how many agents carry out each decisions, like a poll for instance. Balanced union consistency is a stronger requirement than union consistency, but is weaker than argmax and posterior probabilities. The advantage of balanced union consistency is that it put some structures on the decision made on the basis of non-disjoint events. The problem is that we have no decision theory foundations for this decision rule. It may be a direction for future research. The main problem with exhaustiveness is that it is difficult to compare with McKelvey and Page's stochastic regularity.

In Chapter 4, we provided an answer to the second question addressed in the literature,

in the case where the communication protocol may not be public. We showed that if agents have the same utility function and the same prior probability, then communication of the action that maximizes their expected utility eventually leads to consensus in any fair protocol. The advantage of this condition is that it applies to any action spaces. The inconvenient is that we have to make a no-indifference assumption, imposing that the set of maximizing actions is always a singleton.

In Chapter 5, we departed slightly from the two basic questions addressed in the literature. In Parikh and Krasucki [1990]'s setting, the outcome of the communication process depends on the communication protocol, in terms of consensus decision as well as of the amount of information that is learned by the agents during the process. We asked the following question. What happens if agents communicate so as to learn information? We showed that it may be the case that some agents disagree about the protocol to use for communicating, and that this disagreement is common knowledge to all. However, we showed that in this case, the consensus decision does not depend on the protocol. A first limit to this result is that it applies only in decision settings, or in particular game settings such as zero-sum games. A second limit is that preferences over protocols are not complete. Therefore, the situation of common knowledge that two agents disagree about two protocols might be quite unlikely. Finally, we lack a convincing economic interpretation for this result.

In Chapter 6, we examined how the conditions we introduced in Chapters 3 and 4, *i.e.* balanced union consistency and argmax, relate to convexity, weak convexity and union consistency. We showed that there is no inclusion relation between the sets of argmax and convex or weakly convex functions, and no inclusion relation between the sets of balanced union consistent and convex or weakly convex functions. Finally, we showed that, among

the five conditions we identified, convexity and argmax are the only one ensuring that communication of individual decisions leads to consensus in any fair protocol, for more than three agents. Balanced union consistency is the only one ensuring that common knowledge of an exhaustive statistic of individual decisions implies a consensus, for more than three agents.

Bibliography

- [1] Aumann R. J., [1976], Agreeing to Disagree, *The Annals Of Statistics*, **4**, 1236-1239.
- [2] Aumann R., [1987], Correlated equilibrium as an expression of Bayesian rationality, *Econometrica*, **55** (1), 1-18.
- [3] Aumann R., [1999], Interactive epistemology I: knowledge, *International Journal of Game Theory*, **28**, 269-300.
- [4] Aumann R., Brandenburger A., [1995], Epistemic conditions for Nash equilibrium, *Econometrica*, **63** (5), 1161-1180.
- [5] Bacharach M., [1985], Some extensions of a claim of Aumann in an axiomatic model of knowledge, *Journal of Economic Theory*, **37**, 167-190.
- [6] Bassan B., Scarsini M., Zamir S., [1997], "I don't want to know!" Can it be rational?, Discussion Paper #158, The Hebrew University of Jerusalem.
- [7] Bergin J., Brandenburger A., [1990], A simple characterization of stochastically monotone functions, *Econometrica*, **58**, 1241-1243.
- [8] Blackwell D., [1953], Equivalent comparison of experiments, *Annals of Mathematical Statistics*, **24**, 265-272.

-
- [9] Board O., [2004], Dynamic interactive epistemology, *Games and Economic Behavior*, **49**, 49-80.
- [10] Bonanno G., Nehring K., [1997], Agreeing to disagree: a survey, University of California, mimeo.
- [11] Bonanno G., Nehring K., [1999], How to make sense of the common prior assumption under incomplete information, *International Journal of Game Theory*, **28**, 409-434.
- [12] Bonanno G., [2004], A simple modal logic for belief revision, University of California Postprint 1169.
- [13] Cave J., [1983], Learning to agree, *Economics Letters*, **12**, 147-152.
- [14] Chwe M., [2001], *Rational Ritual, Culture, Coordination and Common Knowledge*, Princeton University Press.
- [15] Dekel E., Gul F., [1997], Rationality and knowledge, in *Advances in Economics and Econometrics. Theory and Applications, Volume I*, Kreps D. M. and Wallis K. F. eds, Cambridge University Press.
- [16] Fagin R., Geanakoplos J., Halpern J., Vardi M., [1992], *The expressive power of the hierarchical approach to modeling knowledge and common knowledge*, in M. Vardi, ed., Fourth Symposium on Theoretical Aspects of Reasoning about Knowledge, Los Altos: Morgan Kaufmann Publishers.
- [17] Fagin R., Geanakoplos J., Halpern J., Vardi M., [1999], The hierarchical approach to modeling knowledge and common knowledge, *International Journal of Game Theory*, **28**, 331-365.

-
- [18] Fagin R., Halpern J., Vardi M., [1991], A model-theoretic analysis of knowledge, *Journal of the ACM*, **91**(2), 382-428.
- [19] Fagin R., Halpern J., Moses Y., Vardi M., [1995], *Reasoning About Knowledge*, MIT Press, Cambridge, Massachusetts.
- [20] Geanakoplos J., [1994], Common knowledge, in *Handbook of Game Theory*, vol 2, ed. R. Aumann and S. Hart, New-York.
- [21] Geanakoplos J., Polemarchakis H., [1982], We can't disagree forever, *Journal of Economic Theory*, **28**, 192-200.
- [22] Gossner O., [2000], Comparison of information structures, *Games and Economic Behavior*, **30**, 44-63.
- [23] Gul F. [1998], A comment on Aumann's Bayesian view, *Econometrica*, **66**, 923-927.
- [24] Halpern J., [2001], Alternative semantics for unawareness, *Games and Economic Behavior*, **37**, 321-339.
- [25] Harsanyi J., [1967], Games with incomplete information played by "Bayesian" players, I: The basic model, *Management Science*, **14**, 159-182.
- [26] Heifetz A., [1996], Comment on consensus without common knowledge, *Journal of Economic Theory*, **70**, 273-277.
- [27] Heifetz A., Meier M., Schipper B., [2006], Interactive unawareness, *Journal of Economic Theory*, forthcoming.
- [28] Hintikka J., [1962], *Knowledge and Belief*, Cornell University Press, Ithaca, New-York.
- [29] Hugues G., Cresswell M., [1968], *An Introduction to Modal Logic*, Methuen, London.

-
- [30] Hume D., [1740], *A Treatise of Human Nature*, Clarendon Press, Oxford.
- [31] Jehl D., [1996], Egypt adding corn to bread: an explosive mix? *New York Times*, November 27, p A4.
- [32] Kamien M., Tauman Y., Zamir S., [1990], On the value of information in a strategic conflict, *Games and Economic Behavior*, **2**, 129-153.
- [33] Koessler F., [2001], Common knowledge and consensus with noisy communication, *Mathematical Social Sciences*, **42**, 139-159.
- [34] Krasucki P., [1996], Protocols forcing consensus, *Journal of Economic Theory*, **70**, 266-272.
- [35] Leher E., Rosenberg D., [2003], Information and its Value in Zero-Sum Repeated Games, *mimeo*.
- [36] Lewis D., [1969], *Convention: A philosophical study*, Harvard University Press, Cambridge, Massachusetts.
- [37] Li J., [2006], Information structures with unawareness, University of Pennsylvania, *mimeo*.
- [38] Lipman B., [1995], Approximately common priors, *mimeo*, University of Western Ontario.
- [39] Lismont L., Mongin P., [1994], On the logic of common belief and common knowledge, *Theory and Decision*, **37**, 75-106.
- [40] Littlewood J. E., [1953], *Mathematical Miscellany*, ed. B. Bollobas.

-
- [41] McKelvey R., Page T., [1986], Common knowledge, consensus and aggregate information, *Econometrica*, **54**, 109-127.
- [42] Mertens J.F., Zamir S., [1985], Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory*, **14**, 17-24.
- [43] Milgrom P., [1981], An axiomatic characterization of common knowledge, *Econometrica*, **49**, 219-222.
- [44] Milgrom P., Stockey N., [1982], Information, trade and common knowledge, *Journal of Economic Theory*, **26**, 17-27.
- [45] Morris S., [1995], The common prior assumption, *Economics and Philosophy*, **11**, 227-253.
- [46] Moses Y., Nachum G., [1990], Agreeing to Disagree after all, in R. Parikh (Ed), *Theoretical Aspects of Reasoning about Knowledge: Proc. Third Conference*, Morgan Kaufman.
- [47] Nielsen L., [1995], Common knowledge of a multivariate aggregate statistic, *International Economic Review*, **36**, 207-216.
- [48] Nielsen L., Brandenburger A., Geanakoplos J., McKelvey R., Page T., [1990], Common knowledge of an aggregate of expectations, *Econometrica*, **58**, 1235-1239.
- [49] Parikh R., Krasucki P., [1990], Communication, consensus and knowledge, *Journal of Economic Theory*, **52**, 178-189.
- [50] Rabin M., [1998], Economics and psychologies, *Journal of Economic Literature*, **36**, 11-46.

- [51] Rubinstein A., [1998], *Modeling Bounded Rationality*, MIT Press, Cambridge, Massachusetts.
- [52] Samuelson L., [2004], Modeling knowledge in economic analysis, *Journal of Economic Literature*, **42**, 367-403.
- [53] Savage L., [1954], *The foundations of statistics*, New York: Wiley.
- [54] Schelling T., [1960], *The strategy of conflict*, Harvard University Press, Cambridge, Massachusetts.
- [55] Sebenius J., Geanakoplos J., [1983], Don't bet on it: contingent agreements with asymmetric information, *Journal of the American Statistical Association*, **78**, 424-426.
- [56] Shapley L. S., [1967], On balanced sets and cores, *Naval Research Logistics Quarterly*, **14**, 453-460.
- [57] Weyers S., [1992], Three results on communication, information and common knowledge, CORE Discussion Paper 9228, Université Catholique de Louvain.

Contents

Introduction	1
The notion of common knowledge	2
Motivation and overview of the thesis	7
1 Modelling knowledge	14
1.1 Introduction	14
1.2 Modelling knowledge with Aumann structures	17
1.2.1 Individual knowledge in Aumann structures	17
1.2.2 Properties of knowledge in Aumann structures	20
1.2.3 Common knowledge in Aumann structures	29
1.3 Controversial issues	31
1.3.1 Logical foundations of Aumann structures	34
1.3.2 What do individuals know about the others' knowledge?	37
1.3.3 Do states of the world change over time?	39
1.4 Conclusion	41
2 A review of the <i>Agreeing to Disagree</i> literature	43
2.1 Introduction	43
2.2 Agreeing to disagree 181	47

2.3	Common knowledge of individual decisions negates asymmetric information: the agreement theorems	51
2.3.1	Common knowledge of individual actions	52
2.3.2	Common knowledge of an aggregate of individual decisions	57
2.4	Communication and common knowledge of individual decisions	60
2.4.1	Communication and convergence of beliefs	62
2.4.2	Public communication protocols	65
2.4.3	Non-public protocols	66
2.5	Some criticisms addressed to the Agreeing to Disagree literature	72
2.5.1	Union consistency	72
2.5.2	Like-mindedness	75
2.6	Conclusion	76
3	Consensus and common knowledge of an aggregate of decisions	79
3.1	Introduction	79
3.2	The model	81
3.3	The consensus result	83
3.4	Concluding remarks	88
3.5	Proofs	91
4	Consensus, communication and knowledge: an extension with Bayesian agents	95
4.1	Introduction	95
4.2	Reaching a consensus	97
4.3	Concluding remarks	100

4.4	Proofs	101
5	Communication, consensus and order. Who wants to speak first?	106
5.1	Introduction	106
5.2	Preliminary notions	109
5.3	Who wants to speak first? An agreement theorem.	112
5.4	Concluding remarks	118
5.5	Proof	120
6	Comparative study of some properties of decision rules	122
6.1	Introduction	122
6.2	Relation between union consistency, weak convexity, convexity, balanced union consistency and argmax	123
6.3	Relation between the results	129
6.4	Conclusion	132
	Conclusion	133
	Bibliography	138
	Contents	147