École doctorale IAEM Lorraine

The Dynamics of Incomplete and Inconsistent Information

—

Applications of logic, algebra and coalgebra

La dynamique de l'information incomplète et incohérente

–

Applications de la logique, de l'algèbre et de la coalgèbre

THÈSE

présentée et soutenue publiquement le 5 décembre 2017

pour l'obtention du

Doctorat de l'Université de Lorraine

(mention informatique)

par

Zeinab BAKHTIARINOODEH

Composition du jury

Directeur :	Hans VAN DITMARSCH, Directeur de Recherche CNRS, Université de Lorraine Helle Hvid HANSEN, Maître de Conférences (Co-directeur), Delft University of Technology
Rapporteurs :	Francesco BELARDINELLI, Maître de Conférences, HDR, Université d'Evry Achim JUNG, Professeur, University of Birmingham
Examinateurs :	Miguel COUCEIRO, Professeur, Université de Lorraine Alexander KURZ, Professeur, University of Leicester Stephan MERZ, Directeur de Recherche INRIA, Université de Lorraine Alessandra PALMIGIANO, Maître de Conférences, Delft University of Technology

Laboratoire Lorrain de Recherche en Informatique et ses Applications — UMR 7503

Mis en page avec la classe thesul.

Résumé

Cette thèse est structurée autour de deux axes d'études : (1) développer des logiques épistémiques formalisant la prise en compte de nouvelles données en présence d'informations incomplètes ou incohérentes ; (2) caractériser les notions de bisimulation sur les modèles de ces nouvelles logiques. Les logiques modales utilisées pour formaliser des raisonnements dans le cadre d'informations incomplètes et incohérentes, telle que la logique modale de contingence, sont généralement plus faibles que les logiques modales standards. Nos travaux se basent sur des méthodes logiques, algébriques et co-algébriques.

Dans le Chapitre 2, nous étudions l'influence des connaissances sur le vote stratégique. Nous introduisons une logique pour formaliser la manipulation des procédures de vote dans le cadre d'une connaissance incomplète. Les principaux résultats présentés dans ce chapitre sont: (i) l'introduction d'une logique modale de la connaissance qui tienne compte de l'incertitude d'un électeur concernant les préférences des autres électeurs; (ii) Cela permet de modéliser des scénarios dans lesquels tous les électeurs ont les mêmes préférences, mais peuvent avoir une incertitude différente quant aux préférences des autres électeurs.

Dans le Chapitre 3, nous introduisons une sémantique algébrique pour la logique modale d'action de raffinement. La logique du modèle d'action de raffinement est une extension de la logique du modèle d'action avec des quantificateurs de raffinement. Notre contribution principale est que nous montrons que la logique modale de l'action de raffinement est correcte et complète par rapport à cette sémantique algébrique. Ce chapitre est intéressant car il s'agit d'une première étape vers le développement de contreparties non-classiques de la logique modale de raffinement.

Dans le Chapitre 4, nous introduisons une logique épistémique dynamique pour raisonner sur le changement d'information en présence d'informations incompatibles et incomplètes. La logique que nous présentons dans ce chapitre est une extension de la logique du mode bilattice à quatre valeurs avec des modalités dynamiques. La contribution principale de ce chapitre est une axiomatisation correcte et complète. Ce chapitre contourne les cadres de la logique épistémique dynamique et des logiques modales à plusieurs valeur. Il ouvre la voie à l'étude des fondements mathématiques de la dynamique de la connaissance dans des contextes non-classiques.

Dans le Chapitre 5, nous introduisons la notion de bisimulation pour la logique modale contingence interprété par rapport aux structures des voisinages. La logique de contingence est une extension de la logique propositionnelle avec des modalités de non-contingence. La modalité de contingence peut s'exprimer en fonction de la modalité de nécessité, mais pas l'inverse. Cela rend la logique de contingence moins expressive que la logique modale de base, à la fois sur les modèles Kripke et les modèles des voisinages. Par conséquent, les notions standard de Kripke et de bisimulation des voisinages sont trop fortes pour la logique de contingence. Nous proposons une notion de bisimulation de contingence de voisinage qui correspond à l'expressivité de la logique de contingence. Nos principales contributions dans ce chapitre sont: (i) un théorème de Hennessy-Milner pour la bisimulation de contingence de voisinage; (ii) une caractérisation de la logique de contingence sur les modèles de voisinages comme le fragment invariant de bisimulation de la logique de premier ordre et de la logique modale; (iii) montrant que la logique de contingence possède la propriété d'interpolation Craig.

Dans le Chapitre 6, nous généralisons la notion de bisimulation développée dans le chapitre précédent dans le cadre de la logique modale coalgébrique. Nous introduisons une notion de Λ

-bisimulation pour des logiques modales faiblement expressives et étudions ses propriétés. Le principal résultat technique de ce chapitre est que nous prouvons un théorème Hennessy-Milner pour Λ-bisimulations.

Mots-clés: informations, incomplètes, incohérentes, logique modal, Logique épistémique dynamique, vote, manipulation, modèle Kripke, sémantique algébrique, logique modale de raffinement modèle de voisinage, logique de contingence, logique modal coalgébrique, coalgébra, bisimulation.

Contenu Détaillé

Raisonnement sur le changement d'information

Informations statiques: incomplétude

L'information joue un rôle essentiel dans nos vies. Nous recueillons fréquemment des informations sur Internet, sur les réseaux sociaux, dans les livres ou auprès d'experts afin de prendre des décisions et effectuer nos activités quotidiennes. Malgré le large accès à diverses sources d'information, nous nous trouvons souvent dans des situations où nous avons des informations incomplètes pour prendre des décisions. Des informations incomplètes proviennent, par exemple, d'une observations, mesure inexacte et données manquantes. Dans les conversations quotidiennes, la mention de l'incertitude, de l'imprécision ou de l'inconnu reflète une information incomplète sur une situation.

Dans cette thèse, nous travaillons avec deux formes d'informations incomplètes: l'incertitude et l'information manquante. Ci-dessous nous présentons quelques exemples afin de clarifier ce que nous voulons vraiment dire par les termes "incertitude" et "information manquante".

Example 1 Notre premier exemple est emprunté à [50]. Considérons deux joueurs, Anne et Bill, et un jeu composé de trois cartes différentes: *Coeur*, *Pique* et *Trèfle*. Chaque joueur pioche une carte et ne regarde que sa propre carte. Supposons qu'Anne possède un Coeur, Bill un Pique et que le Trèfle soit placé face cachée sur la table. Anne ne sait pas que Bill détient un Pique, ni que les trèfles sont sur la table. Cependant, elle considère possible que Bill ait un Pique. De même, Bill est incertain quant au fait qu'Anne détienne un Coeur ou un Trèfle. En d'autres termes, Anne et Bill ont une incertitude quant à la répartition réelle des cartes Cette incertitude est une sorte d'information incomplète. ⊣

Example 2 Cathy et sa chef Amy organisent un atelier sur le "Raisonnement sous incertitude". Amy a envoyé un email à Cathy et lui a demandé d'aller vérifier si la salle de conférence était disponible à la date prévue pour l'atelier. Cependant, Cathy, en raison de son emploi du temps chargé, a oublié de vérifier la disponibilité de la salle. Après une semaine, Amy passe dans le bureau de Cathy et lui demande: "Savez-vous si la salle de conférence est disponible pour l'atelier?". Malheureusement, très embarrassée, Cathy doit avouer à sa chef qu'elle a oublié dénvoyer l'email et qu'elle ne sait pas si la salle est disponible. En d'autres termes, elle est incertaine quant à la disponibilité de la salle de conférence à la date prévue pour l'atelier. ⊣

Dans les exemples ci-dessus, la distribution des cartes et la disponibilité de la salle de conférence sont des faits dont la véracité ou la ... sont fixés. L'incertitude porte sur les connaissances des agents. En effet, chaque agent sait quelles sont les options possibles, ils sont juste incertains quant aux choses qui sont vraies ou fausses. Etudions maintenant des exemples un peu différents.

Example 3 Supposons qu'Alice a décidé de planter des tulipes dans son jardin. Elle va au magasin et achète une boîte de bulbes de tulipes. Quand elle arrive à la maison, elle remarque que l'étiquette de la boîte est *manquant*. Ainsi, elle ne sait pas quelle est la couleur des tulipes, et elle ne sait même pas si ce sont effectivement des bulbes de tulipes. De fait, Alice *manque d'informations* pour connaître la couleur et le type de bulbes. Elle devra attendre que les fleures poussent pour connaître le type de bulbes qu'elle a acheté. ⊣

Example 4 Une entreprise ouvre un nouveau poste pour un concepteur de sites Web et crée une base de données composée des noms, des âges et des diplômes des candidats. Ceci est un fragment de cette base de données:

Name	Age	Degree
Steve Cooper		M.Sc.
Mary Lane	27	M.Sc.
John Green	25	B.Sc.

La source responsable d'alimenter la la base de données peut ne pas donner de valeurs pour certains attributs des données. Par exemple, dans la table, nous voyons que la valeur de l'attribut 'age ' pour l'enregistrement relatif à Steve Cooper est manquante. Le comité de sélection ne considère que les candidats âgés d'au plus 28 ans. Ainsi, ils interrogent la base de données pour lister tous les candidats qui ont au maximum 28 ans. Lorsque la base de données répond à cette requête, le nom de Steve Cooper ne figure pas dans la liste, car il n'y a aucune information sur son âge. Le comité de sélection ne considère donc pas sa demande. ⊣

Les informations incomplètes dans les exemples 3 et 4 prennent la forme d'informations manquantes. L'exemple suivant illustre le rôle de l'incertitude dans la théorie du vote .

Example 5 Trois amis Leila, Mona et Sunil veulent voir un film samedi soir. Le cinéma projette Wonder Woman, Life et Logan. Leurs préférences sont (le plus préféré est sur le dessus):

Leila	Mona	Sunil
Wonder Woman	Wonder Woman	Life
Life	Life	Logan
Logan	Logan	Wonder Woman

Cependant, il y a une complication supplémentaire: Leila est incertaine au sujet des préférences de Mona et considère également que Mona préfère Logan over Life et Life over Wonder Woman. Ils vont maintenant voter sur quel film ils iront voir. Lancer un vote signifie déclarer une préférence entière. La règle de vote est la suivante: s'il y a une majorité pour un film préféré, alors ce film gagne, sinon (si les votes sont à égalité) Logan gagne. Sunil jette d'abord son vote et déclare sa vraie préférence (ce vote ne peut plus être changé). Maintenant, Leila et Mona doivent voter simultanément. Que devrait faire Mona? Si Leila et Mona déclarent toutes deux leurs vraies préférences, Wonder Woman gagne et elles sont toutes deux heureuses. Toutefois, Leila considère également possible que Logan soit le film préféré de Mona. Si tel est le cas, et elles déclarent tous les deux leurs véritables préférences, les votes sont à égalité et Logan gagne, le film le moins préféré de Leila. Leila veut éviter la possibilité de ce résultat désagréable, et donc elle décide de déclarer: Je préfère Life sur Wonder Woman et Wonder Woman sur Logan. Maintenant, Life gagne. C'est mieux pour Leila que le résultat Logan. Malheureusement, Mona ne préfère pas Logan à Wonder Woman, mais a la même préférence que Leila. Dans ce cas, le vote alternatif de Leila fait gagner Life, ce qui est pire pour elle que si elle avait voté selon sa vraie préférence: l'incertitude gache son vote! ⊣

Dans l'exemple ci-dessus, Leila n'a pas voté sincèrement, elle a plutôt voté pour Life comme moyen d'obtenir un résultat plus préférable que ce à quoi elle aurait pu s'attendre en votant sincèrement. Dans la théorie du vote, un tel vote s'appelle une manipulation [38, 32]. L'exemple ci-dessus montre que si les électeurs connaissent les préférences des uns et des autres par rapport aux candidats, cela peut affecter le résultat.

Informations statiques: incohérence

Comme nous l'avons mentionné au début, nous obtenons nos informations de plusieurs sources. Les informations que nous obtenons de ces sources peuvent être incohérentes.

Example 6 Imaginez un robot conçu pour sortir d'un labyrinthe. Le robot a deux capteurs qui l'aident à se déplacer le long d'une surface plane dans une direction libre, sans obstacles. Supposons une situation dans laquelle l'un des capteurs détecte un obstacle et l'autre ne détecte rien. En d'autres termes, le robot doit faire face à des informations contradictoires sur lesquelles il doit raisonner pour déterminer dans quelle direction il doit se déplacer. ⊣

Example 7 Reconsidérons l' Exemple 4. Les tableaux suivants représentent maintenant les données des candidats qui ont été fournies par deux sources différentes.

Name	Age	Degree
Steve Cooper	26	M.Sc.
Mary Lane	27	M.Sc.
John Green	25	B.Sc.

Name	Age	Degree
Steve Cooper	26	M.Sc.
Mary Lane	27	Ph.D.
John Green	25	B.Sc.

Les deux sources donnent des informations incohérentes sur le diplôme de Mary Lane. La déclaration "Mary Lane a un doctorat" est à la fois vraie et fausse, car il existe une preuve (Source 2) qu'elle a un doctorat et il y a aussi une preuve (Source 1) qu'elle n'a pas de doctorat. ⊣

Comme l'illustrent les exemples, l'incertitude et l'incohérence sont des caractéristiques communes de l'information. Dans cette thèse, nous nous intéressons à raisonner sur l'information et à analyser les situations dans lesquelles cette information est fournie. Les exemples que nous avons présentés jusqu'à présent décrivent des situations faciles à analyser. Par exemple, dans l'exemple 1, on peut facilement comprendre ce que Anne sait, ou à propos de quoi Bill est incertain. Cependant, les situations peuvent être plus complexes. Par exemple, dans l'exemple 1, il peut y avoir plus de deux joueurs et dans l'exemple 6 il peut y avoir plus d'un robot. Raisonner sur des informations devient donc complexe et exige une analyse formelle plus minutieuse. Nous utilisons la logique pour une telle analyse. Nous passons brièvement en revue les logiques pré-existantes sur lesquelles cette thèse se base. Nous discutons de deux approches, l'une basée sur la logique classique (à deux valeurs) et l'autre sur la logique à quatre valeurs. Dans ce qui suit, nous définissons *knowledge* comme une information vraie et la considérons par rapport à un *agent* (être humain, ordinateur) qui a une certaine perspective sur le monde.

L'analyse du raisonnement sur la connaissance en utilisant la logique formelle remonte à Von Wright [155] et à l'ouvrage fondateur de Hintikka [94], qui a introduit une variante de la logique modale formalisant la notion de connaissance.

Hintikka a fourni une interprétation sémantique de la logique modale de la connaissance en termes de sémantique de Kripke *mondes possibles* [36, 101, 25].L'idée sous-jacente à cette approche est que la connaissance d'un agent peut être caractérisée comme un ensemble de mondes qu'elle considère comme *indiscernables.* Les mondes indiscernables s'appellent les *mondes possibles.* Un agent sait que quelque chose est le cas si et seulement si c'est le cas dans tous les mondes que l'agent considère possible [49]. Cette approche fournit un moyen naturel de modéliser l'incertitude des agents: un agent est incertain de quelque chose si et seulement si c'est le cas dans certains mondes que l'agent considère comme possible et ce n'est pas le cas dans certains mondes que l'agent considère comme possibles. Par exemple, dans l'Exemple 5, on peut modéliser l'incertitude de Leila à propos des préférences de Mona comme suit. Leila considère deux mondes possibles: dans un monde, Mona préfère Wonder Woman sur Life, et Life sur Logan, et dans l'autre monde, elle préfère Logan sur Life, et Life sur Wonder Woman.

La sémantique des mondes possible fournit une manière intuitive et mathématiquement élégante de représenter les connaissances de Mona et Leila et de raisonner sur ces connaissances. Cependant, ce formalisme a certains défauts considérables de notre point de vue: (i) les agents sont *logiquement* [154], c'est-à-dire qu'un agent connaît toutes les conséquences logiques de ses connaissances; (ii) les agents ne peuvent pas avoir d' informations contradictoires sans tout savoir parce que d'une contradiction tout peut être déduit. Ces comportements ne sont pas réalistes et ne sont pas souhaitables lors de la modélisation d' agents limités en terme de ressources [55]. Les agents limités en terme de ressources sont limités en puissance déductive, en durée de raisonnement et en mémoire.

Une logique alternative qui permet aux agents de détenir des connaissances incohérentes sans tout savoir est la *bilattice logic* d'Arieli et Avron [5], qui elle-même est basée sur *la logique à quatre valeurs* de Belnap [15, 16]. dans la logique des En logique bilattice, les propositions peuvent avoir, en plus des valeurs de vérité *true* et *false*, deux autres valeurs, à savoir: (i) *true et false* pour gérer l'incohérence - cela correspond à la situation où plusieurs sources assignent une valeur de vérité différente à une phrase, et (ii) *ni vrai ni faux*, par manque d'information. Il est intéressant de noter que la valeur *à la fois vrai et faux* de la logique de Belnap remonte à la proposition de Łukasiewicz [113] pour aborder le *futur problème de contingence* dans la logique aristotélicienne, qui dit : la vérité des événements futurs ne peut être déterminée dans le présent à moins que nous ayons des informations complètes sur le futur [51]. En ajoutant des modalités à la logique d'Arieli et d'Avron, on peut raisonner sur différentes notions telles que la connaissance, la croyance et le temps. Jung et Rivieccio [98] ont développé une telle extension de la logique des bi-treillis de [5].

Les approches que nous avons mentionnées jusqu'à présent ne décrivent que des états d'information et ne modélisent pas les l'accès à de nouvelles informations. Dans la section suivante, nous discuterons de la dynamique de la connaissance.

Dynamique de l'information

Les connaissances d'un agent peuvent changer en réponse à *un évènement informatif.* Les exemples suivants illustrent que *des* peuvent conduire à des changements dans les connaissances des agents.

Example 8 Reconsidérons l'exemple 1 et supposons qu' Anne triche sans que Bill s'en aperçoive et regarde la carte qui est sur la table. Dans ce cas, Anne apprend la distribution des cartes, tandis que Bill ne sait pas que Anne la connaît . Dans un autre scénario, supposons qu'Anne marche vers la table, prenne la carte face cachée et la regarde sans la montrer à Bill, mais Bill remarque qu'Anne regarde la carte sur la table. Comme précédemment Anne apprend la répartition des cartes et Bill ne connaît toujours pas la distribution des cartes. La différence entre ces deux scénarios réside dans ce que Bill sait des connaissances d'Anne sur les cartes. Dans le premier scénario, Bill ne sait pas qu'Anne connaît la distribution des cartes, alors que dans le deuxième scénario, Bill sait que Anne connaît la distribution des cartes, bien qu'il reste incertain quant à la distribution des cartes lui-même. ⊣

Example 9 Reconsidérons l' Exemple 5, et supposons que Mona, qui dit toujours la vérité, dit à Leila "Je préfère Wonder Woman aux deux autres films, et vous ne le savez pas." En conséquence, Leila sait maintenant que Wonder Woman est leur film préféré, et elle n'est plus ignorante de ce fait. Par conséquent, elle n'est plus incitée à manipuler le vote car son film préféré sera choisi si elle vote sincèrement. ⊣

Il y a deux points dans l'exemple ci-dessus que nous aimerions souligner. Premièrement, l'état d'information de Leila a changé en raison d'une annonce faite par Mona. Cette annonce réduit l'incertitude de Leila et a donc un impact sur son désir de manipuler le vote. Les annonces [135] sont des actions relativement simples qui fournissent aux agents de nouvelles informations susceptibles de résoudre l'ignorance d'un agent à propos d'une situation. Deuxièmement, Mona a annoncé quelque chose qui devient faux après l'annonce. Avant l'annonce, Leila ne connaissait pas les préférences de Mona, mais maintenant qu'elle sait, la déclaration n'est plus vraie. Motivé par ces points, l'un des objectifs de la thèse est d'étudier les aspects théoriques du vote et de la manipulation, avec les *logiques épistémiques dynamiques*.

Logiques épistémiques dynamiques (DEL) [49, 11] sont une famille de logiques pour raisonner sur le changement des connaissances. Ielles étendent la logique modale de la connaissance avec des opérateurs dynamiques.

Une des questions qui nous intéresse avec DEL est : comment la connaissance peut-elle être obtenue par un agent? Cette question relève du sujet *" quantification sur le changement d'information"* [45, 46, 86, 84, 31]. Une logique qui étudie la quantification du changement d'information est appelée *refinement modal logic*. La logique modale de raffinement (RML), introduite par Bozzelli et al. [31], est une extension de la logique modale avec des opérateurs qui quantifient sur tous les *raffinements* d'un modèle. Un raffinement correspond intuitivement au résultat d'une action qui peut changer les états de connaissance des agents.

DEL est un cadre bien adapté pour décrire la dynamique de la connaissance, bien que comme mentionné dans la section 1.1.1, un inconvénient de cette approche est que cela ne tient pas compte de situations où les agents doivent faire face à des connaissances inconsistantes (Exemples 6, 7) ou incomplètes (4). Ces situations sont mieux analysées avec des logiques *non classiques* telles que la logique *à quatre valeurs*. Elle soulève la question de recherche suivante: (i) comment formaliser le changement d'information et quantifier le changement d'information en présence d'informations incomplètes et incohérentes? Pour répondre à cette question nous nous basons sur la théorie de la dualité et appliquons une méthode développée dans [114, 105] et utilisée dans [138] à la logique modale de raffinement et à la logique des bi-treillis développée dans [97].

Expressivité et bisimulation

Le sujet des chapitres 5 et 6 de cette dissertation est de trouver une notion adéquate de *bisimulation* pour les logiques modales *faiblement expressives*.

Bisimulation de Kripke et expressivité de la logique modale de base

Les langages modaux sont utilisés pour exprimer les propriétés des modèles. Différents langages modaux qui sont interprétés sur la même classe de modèles peuvent différer dans leur puissance expressive. L'expressivité est une mesure du pouvoir qu'a un langage modal de distinguer différents modèles. Donnons un exemple qui est emprunté à [26] avec de légères modifications. Considérez les modèles montrés ci-dessous.

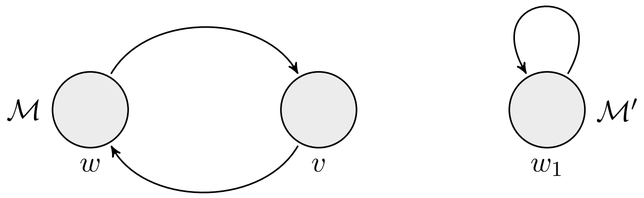

Le modèle $\mathcal{M}'$ est réflexif, alors que le modèle $\mathcal{M}$ n'est pas réflexif, et comme la réflexivité peut être exprimée comme une formule de premier ordre [25], cela signifie que les modèles $\mathcal{M}$ et $\mathcal{M}'$ se distinguent par une formule de premier ordre. Cependant, il n'y a pas de formule dans le langage modal de base qui puisse distinguer $\mathcal{M}$ et $\mathcal{M}'$, c'est-à-dire qu'ils sont *modalement équivalent.* Concernant cet exemple, on peut se demander pourquoi ces modèles sont indistingables ? Ou plus généralement, on peut se demander quand deux modèles distincts sont modalement équivalents. Un outil standard utilisé pour répondre à ces questions sont les bisimulations [17]. Il est bien connu que les formules modales sont invariantes par des bisimulations, ce qui signifie que si deux états dans un modèle de Kripke sont bisimilaires (c'est-à-dire liés par une bisimulation), alors ils sont modalement équivalents. L'inverse est généralement faux (voir, par exemple, [25, Exemple 2.23]), mais c'est cependant vrai sur la classe des *modèles Kripke à image-finis.* Ce résultat est appelé le théorème de Hennessy-Milner pour la logique modale de base [93].

Un autre lien important entre la bisimulation et l'expressivité de la logique modale a été établi par le théorème de caractérisation de Van Benthem [17]. Ce théorème dit que toute propriété invariante par bisimulation sur les modèles de Kripke qui peut être définie dans la logique du premier ordre, est également définissable dans la logique modale de base.

Logiques modales faiblement expressives

Nous disons qu'un langage modal est faiblement expressif, si tout ce que nous pouvons dire dans ce langage peut s'exprimer dans la logique modale de base, mais pas vice versa. Un exemple d'un tel langage modal est donné par la logique de contingence [62] dans laquelle les modalités désignent la (non) contingence des propositions. Une proposition est non-contingente si elle est nécessairement vraie ou nécessairement fausse, sinon elle est contingente. Dans un contexte épistémique, "la proposition est non-contingente" devient "l'agent sait si la proposition est vraie" et "la proposition est contingente" devient "l'agent est incertain si la proposition est vraie". La notion de "savoir si" peut être exprimée en termes de "savoir cela": un agent sait si une proposition est vraie, si il sait qu'elle est vraie ou si il sait qu'elle est fausse. Aussi, "savoir que" peut être exprimé comme "savoir si": un agent sait qu'une proposition est vraie, si la proposition est vraie et qu'il sait si la proposition est vraie. Cependant, en l'absence de la propriété de la connaissance que les propositions connues sont vraies, la notion de nécessité (sachant que) ne peut pas nécessairement être définie en termes de la notion de non-contingence (savoir si). La logique de contingence est donc moins expressive que la logique modale de base. Deux sémantiques différentes sont proposées pour la logique de contingence, l'une basée sur les modèles de Kripke [62, 63], et l'autre basée sur les modèles de voisinages [61]. Ce dernier est motivé par le fait que la non-contingence en tant qu'opérateur modal n'est pas monotone, et donc la logique de contingence n'est pas une logique modale normale. Les modèles de voisinages [37, 121] sont une généralisation des modèles de Kripke, et ils sont devenus l'outil sémantique standard pour raisonner sur les logiques modales non normales.

Énoncé du problème Comme mentionné ci-dessus, le théorème de Hennessy-Milner et le théorème de caractérisation de Van Benthem montrent que la bisimilarité de Kripke correspond au pouvoir expressif du langage modal de base sur les modèles de Kripke. Ces résultats ont été généralisés aux modèles de voisinages et à la bisimilarité de voisinages dans [89]. Cependant, pour des logiques faiblement expressives, telles que la logique de contingence, les notions de bisimilarité de Kripke et de bisimilarité de voisinages sont trop fortes, ce qui signifie que deux états qui satisfont les mêmes formules peuvent ne pas être Kripke/voisinage bisimilaires. Pour remédier à cette situation pour la logique de contingence, Fan et al. [62] proposent une notion de bisimulation de contingence sur des modèles de Kripke, et ils montrent un théorème de Hennessy-Milner et un théorème de caractérisation. Cependant, leur définition est trouvée de manière ad hoc, et il n'est pas clair comment généraliser les bisimulations de contingence aux modèles de voisinages. Les questions suivantes se posent donc: Quelle est la bonne notion de bisimulation de contingence de voisinages? Et plus généralement, avec une logique faiblement expressive, peut-on systématiquement définir une notion de bisimulation?

Méthodologie et travail existant Pour aborder les questions ci-dessus, nous passons à un niveau d'abstraction plus élevé et utilisons les coalgèbres et les logiques coalgébriques *Les coalgèbres* [141] peuvent être vues comme une abstraction de systèmes basés sur l'état tels que les systèmes de transitions étiquetés, les cadres de Kripke et les structures de voisinage [82]. Informellement, une coalgebra est un ensemble de mondes avec une carte de transition dont le type est paramétrique dans le choix du foncteur. La théorie des coalgebres est étroitement liée à la logique modale de deux manières: Premièrement, les modèles de logique modale tels que les cadres de Kripke et les structures de voisinage peuvent être représentés comme des coalgebres. Ainsi, les coalgebras généralisent les modèles traditionnels de la logique modale. Deuxièmement, il est montré que la logique modale est un outil adéquat pour raisonner sur les coalgebres, dans le même sens que la logique équationnelle est la logique de base des algèbres [104]. En effet, les chercheurs ont introduit une théorie générale connue sous le nom de *logique modale coalgebrique* [124, 131, 132] comme cadre général dans quelles logiques modales pour différents types de structures peuvent être développées dans un cadre uniforme. La logique modale coalgebrique développée par Pattinson dans [131] utilise ce que l'on appelle des élévations de prédicats pour définir un langage modal. Informellement, une levée de prédicat peut être vue comme une généralisation des modalités de nécessité (possibilité) de la logique modale de base. Logique modale Coalgebraic est livré avec des méthodes génériques et outils pour prouver la solidité et l'exhaustivité [131, 145], la décidabilité [144], l'expressivité [132, 143], et développer la théorie de la correspondance [111]. Notre motivation à utiliser un cadre coalgebraic est double. Premièrement, en raison de la paramétricité du type mentionnée ci-dessus, la houille bitumineuse permet de développer des définitions et des résultats uniformes pour différents types de structures. Ces résultats peuvent ensuite être instanciés pour des classes de structures concrètes et leurs logiques modales. Deuxièmement, la théorie des coalgebras s'accompagne de notions générales d'équivalence entre états, à savoir les bisimulations coalgebriques [1, 141] et l'équivalence comportementale [104], ainsi que les résultats généraux sur l'invariance de bisimulation et expressivité [132, 143].

La logique modale coalgebrique, telle qu'elle est consideree dans [131], est invariante par une equivalence de comportement, mais l'inverse ne tient pas en general. Pattinson dans [132] a proposé une condition sur les élévations de prédicats sous laquelle la logique modale binébrique est *expressive*, signifiant que si deux états satisfont les mêmes formules, alors ils sont comportementaux équivalents. Le travail existant dans la logique modale coalgebrique s'est concentré sur l'identification des conditions qui assurent qu'un langage modal est expressif. Le but de cette thèse est de changer les choses et commencer à partir d'un faible expressive coalgebraic logique modale, et trouver une notion de bisimulation qui correspond à son expressivité. Nos résultats donc aussi contribuer à la théorie générale de modal coalgebrique logique.

Organisation et contributions

Nous concluons l'introduction en donnant un aperçu des chapitres de cette thèse, leur thèmes de recherche et un bref énoncé de leurs contributions.

Dans le Chapitre 2, nous étudions l'influence des connaissances sur le vote stratégique. Nous introduisons une logique pour formaliser la manipulation des procédures de vote dans le cadre d'une connaissance incomplète. Les principaux résultats présentés dans ce chapitre sont: (i) l'introduction d'une logique modale de la connaissance qui tienne compte de l'incertitude d'un électeur concernant les préférences des autres électeurs; (ii) Cela permet de modéliser des scénarios dans lesquels tous les électeurs ont les mêmes préférences, mais peuvent avoir une incertitude différente quant aux préférences des autres électeurs.

Dans le Chapitre 3, nous introduisons une sémantique algébrique pour la logique modale d'action de raffinement. La logique du modèle d'action de raffinement est une extension de la logique du modèle d'action avec des quantificateurs de raffinement. Notre contribution principale est que nous montrons que la logique modale de l'action de raffinement est correcte et complète par rapport à cette sémantique algébrique. Ce chapitre est intéressant car il s'agit d'une première étape vers le développement de contreparties non-classiques de la logique modale de raffinement.

Dans le Chapitre 4, nous introduisons une logique épistémique dynamique pour raisonner sur le changement d'information en présence d'informations incompatibles et incomplètes. La logique que nous présentons dans ce chapitre est une extension de la logique du mode bilattice à quatre valeurs avec des modalités dynamiques. La contribution principale de ce chapitre est une axiomatisation correcte et complète. Ce chapitre contourne les cadres de la logique épistémique dynamique et des logiques modales à plusieurs valeur. Il ouvre la voie à l'étude des fondements mathématiques de la dynamique de la connaissance dans des contextes non-classiques.

Dans le Chapitre 5, nous introduisons la notion de bisimulation pour la logique modale contingence interprété par rapport aux structures des voisinages. La logique de contingence est une extension de la logique propositionnelle avec des modalités de non-contingence. La modalité de contingence peut s'exprimer en fonction de la modalité de nécessité, mais pas l'inverse. Cela rend la logique de contingence moins expressive que la logique modale de base, à la fois sur les modèles Kripke et les modèles des voisinages. Par conséquent, les notions standard de Kripke et de bisimulation des voisinages sont trop fortes pour la logique de contingence. Nous proposons une notion de bisimulation de contingence de voisinage qui correspond à l'expressivité de la logique de contingence. Nos principales contributions dans ce chapitre sont: (i) un théorème de Hennessy-Milner pour la bisimulation de contingence de voisinage; (ii) une caractérisation de la logique de contingence sur les modèles de voisinages comme le fragment invariant de bisimulation de la logique de premier ordre et de la logique modale; (iii) montrant que la logique de contingence possède la propriété d'interpolation Craig.

Dans le Chapitre 6, nous généralisons la notion de bisimulation développée dans le chapitre précédent dans le cadre de la logique modale coalgébrique. Nous introduisons une notion de Λ -bisimulation pour des logiques modales faiblement expressives et étudions ses propriétés. Le principal résultat technique de ce chapitre est que nous prouvons un théorème Hennessy-Milner pour Λ-bisimulations.

Abstract

In the present Ph.D. dissertation we investigate reasoning about information change in the presence of incomplete or inconsistent information, and the characterization of notions of bisimulation on models encoding such reasoning patterns. Modal logics for incomplete and inconsistent information are typically weaker than the standard modal logics, such as the modal logic of contingency. We use logical, algebraic and co-algebraic methods to achieve our aims. The dissertation consists of two main parts. The first part focusses on reasoning about information change, and the second part focusses on expressivity and bisimulation. In the following, we give an overview of the contents of this dissertation.

In Chapter 1, we give an introduction to the main topics of the thesis.

Chapter 2 studies the influence of knowledge on strategic voting. We introduce a logical framework to study manipulation of voting procedure under incomplete knowledge. The fact that voters may or may not know each other's preferences will affect their ability to manipulate. Our main contributions in this chapter are: (i) the introduction of a modal logic of knowledge to account for a voter's uncertainty about other voters' preferences; (ii) addressing the question of how a reduction in uncertainty may affect manipulation. The main merit of these contributions is that they enable us to model higher-order knowledge. For example, we can model that voter 1 is uncertain about voter 2's preferences, and voter 1 knows that voter 3 knows voter 2's preferences. This makes it possible to model scenarios wherein all voters have the same preferences, but may have different uncertainty about the preferences of other voters. In Chapter 3, we introduce an algebraic semantics for refinement action modal logic. Refinement action model logic is an extension of action model logic with refinement quantifiers. Our main contribution is that we show that refinement action modal logic is sound and complete with respect to this algebraic semantics. This work is of interest as it a first step towards developing non-classical counterparts of refinement modal logic.

In Chapter 4, we develop dynamic epistemic logic for reasoning about information change in the presence of inconsistent and incomplete information. The logic we introduce in this chapter is an extension of the four-valued bilattice modal logic with dynamic modalities. The main contribution of this chapter is a sound and complete axiomatisation. This work bridges the frameworks of dynamic epistemic logic and many-valued modal logics. It paves the way to study the mathematical foundations of dynamics of knowledge in non-classical settings.

In Chapter 5, we introduce the notion of bisimulation for contingency logic interpreted over neighbourhood models. Contingency logic is an extension of propositional logic with (non-)contingency modalities. The contingency modality can be expressed in terms of the necessity modality, but not the other way around. This makes contingency logic less expressive than basic modal logic, both over Kripke models and neighbourhood models. Hence, the standard notions of Kripke and neighbourhood bisimulation are too strong for contingency logic. We propose a notion of neighbourhood contingency bisimulation that fits the expressivity of contingency logic. Our main contributions in this chapter are: (i) a Hennessy-Milner theorem for neighbourhood contingency bisimulation; (ii) a characterisation of contingency logic over neighbourhood models as the bisimulation invariant fragment of first-order logic and of modal logic; (iii) showing that contingency logic has the Craig interpolation property.

In Chapter 6, we generalise the notion of bisimulation developed in the previous chapter to the framework of coalgebraic modal logic. We introduce a notion of Λ-bisimulation for weakly expressive coalgebraic modal logics, and study its properties. The main technical result of this chapter is that we prove a Hennessy-Milner theorem for Λ-bisimulations.

Keywords: incomplete knowledge, inconsistent knowledge, modal logic, dynamic epistemic logic, voting, manipulation, Kripke model, algebraic semantics, refinement action modal logic, neighbourhood model, contingency logic, coalgebraic modal logic, coalgebra, bisimulation

To Maryam Mirzakhani, a great mathematician
and a wonderful human being who broke a glass ceiling
and inspired many, men and women alike,
and to my love, Ahmadreza.

Acknowledgments

I would like to thank my supervisors, Hans van Ditmarsch and Helle Hvid Hansen: thanks for the incredible support and guidance you offered throughout my PhD. I believe I have been privileged.

I would like to express my sincere gratitude to my thesis committee: Prof. Jung, Prof. Kurz, Prof. Couceiro, Dr. Palmigiano, Dr. Merz, Dr. Belardinelli and Dr. Balbiani for their insightful and constrictive comments.

A huge thanks goes to Sabine Frittella for being an amazing advisor and exceptional friend. It is hard to imagine having made it through this Ph.D. without her companionship. Sabine provided the French translation of the results of my Ph.D. thesis, for which also many thanks.

I would like to thank my nice and supportive office-mates, Aybüke and Sophia. Aybüke has been always a great listener and an incredible friend. Sophia has been always there for me and answered all my questions patiently.

During my stay in Nancy, I have gained nice colleagues, who made my Ph.D. very interesting. Christophe, thanks for all support and chats. Louwe, thanks for the discussions that made our lunch time informative.

I would like to thank my collaborators Umberto Rivieccio and Abdallah Saffidine. I have learned a lot from you and enjoyed working with you. I hope I get the opportunity to collaborate with you again in future.

My friends from around the world have all been an important part of this journey. Shabnam, thank you for more than a decade of camaraderie. Zia and Masoumeh, thanks for listening to my woes.

I am deeply indebted to my master advisor, Dr. Majid Alizadeh, who supported me from the first year of my bachelor studies and encouraged me to start my Ph.D. in France.

A very special word of thanks goes to my family. Thanks to my parents for being the source of my strength and inspiration. They gave me the courage to follow my dreams even when they seemed unattainable. I could not have done this without you. My amazing siblings, Zahra, Mostafa, and Ali, thank you for all the laughs, the love and being by my side even when we were far apart. My amazing niece, Mohaddeseh, thank you for making my life so amazing with all your sweet words. I also would like to express my gratitude to my parents-in-low, thanks for all praying and positive energies.

Last but certainly not least, I would like to express the deepest appreciation to my beloved husband Ahmadreza. You are more than my husband, you are my best friend who always believes in me even when I myself do not. I am so lucky to have you in my life.

Contents

List of Figures

List of Tables

List of Symbols

Ag	set of voters(agents): 13
i, j	voter variables: 13
$\mathcal{C}$	set of candidates: 13
x_i, y_i	candidate variables: 13
$\succ_i$	preference relation for voter i: 13
$O(\mathcal{C})$	set of all linear orders on $\mathcal{C}$: 13
P, P', P''	profile variables: 13
P_i	vote cast by i: 13
F	voting rule function: 13, bifilter: 71
$O(\mathcal{C})^n$	set of all profiles: 13
P_{-i}	profile without P_i: 13
$P(P_{-i}, \succ_i')$	substitution of P_i by $\succ_i'$ in P: 13
$\succ$	tie-breaking mechanism: 13
$\mathcal{M}$	Kripke model: 15, 67
S	set of states: 15, 39
s	state variable: 15
$\sim_i$	indistinguishability relation for voter i: 15
V	profile function: 15, valuation: 39
$[s]_{\sim_i}$	i-equivalence class of s: 15
$\mathfrak{P}$	collection of profiles: 20
$(\mathcal{M}, s)$	knowledge profile pointed Kripke model: 15, 67
$[\succ]_i$	conditional vote for voter i: 23
$((\mathcal{M}\|T), s)$	restricted knowledge model to subdomain T: 28
$((\mathcal{M}\|\varphi), s)$	updated model φ: 28
$\mathcal{L}_{KV}$	language of the logic of knowledge and voting: 31
$\wedge$	conjunction: 31, meet: 45
$\neg$	negation: 31, complement: 45, two-valued negation: 65
$\vee$	disjunction: 31, join: 45
$\rightarrow$	implication: 31, strong implication: 66
$\leftrightarrow$	bi-implication: 31, strong bi-implication: 66
K_i	knowledge modality for agent i: 31
$[\varphi]$	public announcement modality: 32
$\bigvee$	infinite disjunction: 32
$\bigwedge$	infinite conjunction: 33
At	set of propositional variables: 66
$p, q, r, ...$	propositional variables: 39
$\mathcal{L}_\Box$	language of basic modal logic: 39
$\Box$	necessity: 39
$\Diamond$	possibility: 39

$\mathcal{F}$	(four-valued) Kripke frame: 39
R	accessibility relation: 39
$\langle \mathcal{F}, s \rangle$	pointed Kripke frame: 39
$\langle \mathcal{F}, T \rangle$	multi-pointed Kripke frame: 39
$(\mathcal{M}, T)$	multi-pointed Kripke model: 39
$[\![\varphi]\!]_{\mathcal{M}}$	truth set of φ in $\mathcal{M}$: 40
$\mathbf{K}$	basic modal logic: 40
$\models_{\mathbf{K}}$	semantic consequence relation in **K**: 40
∇	cover modality: 40, contingency modality: 104
t	always true proposition: 40, bilattice value for true: 63
f	always false proposition: 40, bilattice value for false: 63
$\leftrightarroweq$	Kripke bisimilarity: 40
$\mathbb{K}$	axiomatisation for **K**: 40
$\vdash_{\mathbb{K}}$	derivability relation in $\mathbb{K}$: 40
$\mathcal{L}_{\Box\alpha}$	language of action model logic: 41
$\mathcal{AM}_{AML}$	set of action models (over $\mathcal{L}_{\Box\alpha}$): 41
α	action model: 41
K	domain of an action model: 41
R_{α}	accessibility relation of an action model: 41
Pre_{α}	precondition of an action model: 41
α_k	pointed action model: 41
α_T	multi-pointed action model: 41
$[\alpha_k]$	necessity action model modality: 41
$\langle \alpha_k \rangle$	possibility action model model quantifier: 41
$[\alpha_T]$	necessity multi-pointed action model modality: 41
$\langle \alpha_T \rangle$	possibility multi-pointed action model modality: 41
$\mathcal{M}_{\alpha}$	product update: 41
$S_{\times}$	domain of product update: 41
$R_{\times}$	accessibility relation of product update: 41
$V_{\times}$	valuation of product update: 41
AML	logic of action model logic: 41
$\models_{\mathbf{AML}}$	semantic consequence relation in **AML**: 41
$\mathbb{AML}$	axiomatisation for **AML**: 41
$\vdash_{\mathbb{AML}}$	derivability relation in $\mathbb{AML}$: 42
$\mathfrak{R}$	refinement: 42
$\mathcal{L}_{\Box\forall}$	language of refinement modal logic: 42
$\forall$	refinement universal quantifier: 42
$\exists$	existential refinement modality: 42
$\models_{\mathbf{RML}}$	semantic consequence relation in **RML**: 42
RML	refinement modal logic: 42
$\mathbb{RML}$	axiomatisation for **RML**: 43
$\vdash_{\mathbb{RML}}$	derivability relation in $\mathbb{RML}$: 43
$\mathcal{L}_{\Box\alpha\forall}$	language of refinement action model logic: 43
$\mathcal{AM}_{RAML}$	set of action models (over $\mathcal{L}_{\Box\alpha\forall}$): 43
RAML	refinement action model logic: 43
$\models_{\mathbf{RAML}}$	semantic consequence relation in **RAML**: 43
$\mathbb{RAML}$	axiomatisation for **RAML**: 43
$\vdash_{\mathbb{RAML}}$	derivability relation in $\mathbb{RAML}$: 43

$\alpha^{\varphi}_{K^{\varphi}}$	multi-pointed action model for φ: 44
$\mathbf{A}$	Boolean algebra: 45
$\mathbf{0}$	bottom element in Boolean algebras: 45
$\mathbf{1}$	top element in Boolean algebras: 45
$\mathbf{BA}$	class of Boolean algebras: 45
$(\mathbf{A}, \leq)$	poset: 45
$\mathcal{P}(X)$	powerset of X: 46
$\cap$	intersection: 46
$\cup$	union: 46
$(-)^c$	complement: 46
$\mathbb{A}$	modal Boolean algebra: 45
$\mathbf{MBA}$	class of modal Boolean algebra: 45
$\cong$	isomorphism: 46
θ	congruence: 46
A/θ	quotient algebra: 46
$\dashv$	adjunction: 47
$\blacklozenge$	left adjoint of $\Box$: 47
$\blacksquare$	right adjoint of $\Diamond$: 47
$\langle \mathbb{A}, \blacksquare \rangle$	tense algebra: 47
$\mathcal{F}^+$	complex algebra: 47
$\mathcal{A}$	algebraic model: 48
$\coprod_\alpha \mathcal{M}$	intermediate model: 48
$\coprod_K S$	domain of intermediate model: 48
$R \times R_\alpha$	accessibility relation in intermediate model: 48
$\coprod_\alpha V$	valuation of intermediate model: 48
a	algebraic action model: 49
R_a	accessibility relation in algebraic action model: 49
Pre_a	precondition in algebraic action model: 49
$\prod_a \mathbb{A}$	intermediate modal Boolean algebra: 49
$\prod_\alpha \mathcal{A}$	intermediate algebraic model: 51
$\prod_a \mathbb{A}$	intermediate algebra: 51
$\equiv_a$	pseudo-quotient relation: 50
$\mathbb{A}_a$	updated algebra: 50
$\Box^a$	necessity modality of the updated algebra: 50
$\Diamond^a$	possibility modality of the updated algebra: 50
q	quotient map: 51
$\mathcal{A}_\alpha$	algebraic updated model: 51
V_a	valuation of updated algebraic model: 51
i'	injection map (adjoint of q): 51
$\mathfrak{A}_{\mathcal{A}}$	refinement algebra: 55
$\mathbb{A}^{\varphi}$	updated algebra for φ: 53
$(b^{\varphi})_{\varphi \in \mathcal{L}_{\Box\forall}}$	elements of the refinement algebra 55
$\top$	true and false: 63
$\bot$	neither true nor false: 63
$\mathbf{4}$	Belnap's four element set: 63
$\leq_t$	truth order: 63
$\leq_k$	knowledge order: 63
FOUR	Belnap bilattice: 63

$\mathcal{L}_\Lambda$	language of coalgebraic modal logic: 133
$(T, ([\![\heartsuit]\!])_{\heartsuit\in\Lambda})$	Λ-structure: 133
$\heartsuit$	modal operator: 133
$[\![\heartsuit]\!]$	predicate lifting: 133
$pb(f, f')$	(weak) pullback of two functions: 130
$\equiv_\Lambda$	modally equivalent over $\mathcal{L}_\Lambda$: 134
$\Box_n$	instantial modal operator: 135
$\sim_{\Lambda_+}$	Λ-bisimilarity between coalgebras: 136
$\sim_{bh}$	behavioural equivalence: 132
$\sim_T$	T-bisimilarity: 133
$\sim_p$	precocongruence: 133
$\sim_\Lambda$	Λ-bisimilarity on coproduct: 136

1
Introduction

> "What one knows is not as much as what one does not know. There is a great variety of things."
>
> Ko Huang

In this Ph.D. dissertation we investigate reasoning about information change in the presence of incomplete or inconsistent information, and the characterisation of notions of bisimulation on models encoding such reasoning patterns. Modal logics for incomplete and inconsistent information are typically weaker than the standard modal logics, such as the modal logic of contingency. We use both logical, algebraic and co-algebraic methods to achieve our aims. The dissertation consists of two main parts. The first part focusses on reasoning about information change, and the second part focusses on expressivity and bisimulation.

1.1 Reasoning about information change

The topic of Chapters 3, 4 and 5 of the dissertation is the study of reasoning about information change in the presence of incomplete or inconsistent information.

1.1.1 Static information: incompleteness

Information plays a vital role in our lives. We frequently acquire information from the internet, social media, books, and experts to make decisions and do our daily activities. Despite the broad access to various information sources, we often find ourselves in situations where we have *incomplete* information to make decisions. Incomplete information arises from, e.g., insufficient observations, inaccurate measurement, and missing data. In daily conversations, the mention of uncertainty, vagueness, or the unknown reflects one's incomplete information about a situation. In this dissertation, we work with two forms of incomplete information: uncertainty and missing information. Let us provide some detailed examples such that it becomes clear what we mean by uncertainty and missing information.

Example 1.1.1 Our first example is borrowed from [50]. Consider two players, Anne and Bill, and a deck consisting of three different cards: *Hearts*, *Spades* and *Clubs*. Each player draws a card and looks only at her/his own card. Suppose Anne holds Hearts, Bill holds Spades, and Clubs is put facedown on the table. Then, Anne does not know that Bill holds Spades and does not know that Clubs is on the table. However, she considers it possible that Bill has Spades. Similarly, Bill is uncertain whether Anne holds Hearts or Clubs. In other words, Anne and Bill are *uncertain* about the actual deal of the cards. This uncertainty is a kind of incomplete information. ⊣

Example 1.1.2 Cathy and her boss Amy organise a workshop on "Reasoning under uncertainty". Amy sent an email to Cathy and asked her to go and check whether the conference room is available on the planned date of the workshop. However, Cathy, due to her busy schedule, forgot to check the availability of the room. After a week, Amy stops by Cathy's office and asks her: "Do you know if the conference room is available for the workshop?". Unfortunately, to her own embarrassment, Cathy must admit to her boss that she forgot the email and she does not know if the room is available. In other words, she has *uncertainty* about the availability of the conference room on the planned date of the workshop. ⊣

In the above examples, the incomplete information takes the form of uncertainty, in which the actual state of affairs is completely determined but we know the possibilities. Let us continue with more examples.

Example 1.1.3 Assume Alice has decided to plant tulips in her garden. She goes to the store and buys a box of tulip bulbs. When she arrives home, she notices that the label of the box is *missing*. Thus, she does not know what the colour of the tulips is, and she even does not know whether they are tulip bulbs. Hence, Alice is *missing information* about the colour and the type of the bulbs and she cannot find out until they grow and are in bloom.

Example 1.1.4 A company opens a new position for a web designer and makes a database consisting of applicants' names, ages, and degrees. This is a fragment of this database:

Name	Age	Degree
Steve Cooper		M.Sc.
Mary Lane	27	M.Sc.
John Green	25	B.Sc.

The source responsible providing information to the database may fail to give values for some attributes of the data. E.g., in the table, we see that the value of the 'age' attribute for the record relating to Steve Cooper is missing. The selection committee only considers applicants who are at most 28 years old. So, they query the database to list all applicants who are at most 28. When the database answers this query, the name of Steve Cooper is not in the list, because there is no evidence on his age. So, the selection committee does not consider his application. ⊣

The incomplete information in Examples 1.1.3 and 1.1.4 takes the form of missing information.

The next example illustrates the role of uncertainty in *voting theory*.

Example 1.1.5 Three friends Leila, Mona, and Sunil want to see a movie on Saturday night. The theatre shows Wonder Woman, Life, and Logan. Their preferences are (most preferred is on top):

Leila	Mona	Sunil
Wonder Woman	Wonder Woman	Life
Life	Life	Logan
Logan	Logan	Wonder Woman

However, there is an additional complication: Leila is uncertain about Mona's movie preferences and also considers it possible that Mona prefers Logan over Life, and Life over Wonder Woman. They will now vote on which movie they will go to. Casting a vote means declaring an entire preference. The voting rule is as follows: if there is a majority for a most preferred movie, then that movie wins, otherwise (if the votes are tied) Logan wins. Sunil casts his vote first and declares his true preference (this vote can no longer be changed). Now Leila and Mona have to cast their votes simultaneously. What should Mona do? If Leila and Mona both declare their true preference, Wonder Woman wins and they are both happy. However, Leila also considers it possible that Logan is Mona's preferred movie. If that is so, and they both declare their true preference, the votes are tied and Logan wins, Leila's least preferred movie. Leila wants to avoid the possibility of that unpleasant outcome, and therefore she decides to declare: I prefer Life over Wonder Woman and Wonder Woman over Logan. Now Life wins. This is better for Leila than the outcome Logan. Unfortunately, Mona does not really prefer Logan over Wonder Woman, but has the same preference as Leila. In this case, Leila's alternative vote makes Life win, which is a worse outcome for her than if she had voted according to her real preference: the uncertainty spoils her vote! ⊣

In the above example, Leila did not cast her sincere vote, instead she voted for Life as a way of obtaining a more preferable outcome than what she would have expected by voting sincerely. In the theory of voting, such a vote is called a *manipulation* [38, 32]. The above example shows that whether voters know eachother's preferences over the candidates may effect the outcome.

1.1.2 Static information: inconsistency

As we mentioned early on, we obtain our information from multiple sources. The information that we obtain from these sources may be inconsistent.

Example 1.1.6 Consider a robot that has been designed to find its way out of a maze. The robot has two sensors that help it to move along a flat surface in a free direction, without obstacles. Assume a situation in which one of the sensors detects an obstacle and the other one detects nothing. In other words, the robot has to cope with contradictory information about which it should reason in order to determine in which direction it should move. ⊣

Example 1.1.7 Reconsider Example 1.1.4. The following tables now represent the information of the applicants that has been provided by two different sources.

Name	Age	Degree
Steve Cooper	26	M.Sc.
Mary Lane	27	M.Sc.
John Green	25	B.Sc.

Name	Age	Degree
Steve Cooper	26	M.Sc.
Mary Lane	27	Ph.D.
John Green	25	B.Sc.

The two sources give inconsistent information on the degree of Mary Lane. The statement "Mary Lane has a Ph.D. degree" is both true and false, because there is an evidence (Source 2) that she has a Ph.D. and there also is an evidence (Source 1) that she does not have a Ph.D. ⊣

As the examples illustrate, uncertainty and inconsistency are common features of information. In this dissertation, we are interested in reasoning about information and analysing the situations in which such information is provided. The examples we presented so far describe situations that can be analysed easily. For instance, in Example 1.1.1, one can easily figure out what Anne knows, or about what Bill is uncertain. However, the situations can be more complex. For example, in Example 1.1.1 there might be more than two players and in Example 1.1.6 there might be more than one robot. Then, reasoning about information can become intricate and demand more careful formal analysis. We use logic for such an analysis. We briefly review the existing work on logics for reasoning about knowledge on which this dissertation builds. We discuss two approaches, one based on classical (two-valued) logic, and one based on four-valued logic. In the following, we define *knowledge* as true information and consider it relative to an *agent* (human-being, computer) who has a certain perspective on the world.

Formal logical analysis of reasoning about knowledge goes back to Von Wright [155] and the seminal work by Hintikka [94], who introduced a variant of modal logic as a formal account of knowledge. Hintikka provided a semantic interpretation of the modal logic of knowledge in terms of a Kripke-style *possible worlds* semantics [36, 101, 25]. The idea behind this approach is that an agent's knowledge can be characterised as a set of worlds that she considers to be *indistinguishable*. The indistinguishable worlds are called the *possible worlds*. An agent knows that something is the case if and only if it is the case in all the worlds that the agent considers possible [49]. This approach provides a natural way to model uncertainty of the agents: an agent is uncertain about something if and only if it is the case in some worlds that the agent considers possible and it is not the case in some worlds that the agent considers possible. For instance, in Example 1.1.5, one can model Leila's uncertainty about Mona's preferences as follows. Leila considers two possible worlds: in one world Mona prefers Wonder Woman over Life, and Life over Logan, and in the other world she prefers Logan over Life, and Life over Wonder Woman. The possible worlds semantics of knowledge provides an intuitive and mathematically elegant way of representing and reasoning about knowledge. However, it has some defects that are important from our perspective: (i) agents are *logical omniscient* [154], that is, an agent knows all logical consequences of her knowledge; (ii) agents cannot hold contradictory information without knowing everything because from a contradiction everything can be inferred. These features are not realistic and not desirable when modelling resource-bounded agents [55]. Resource bounded agents are limited in deductive power, duration of reasoning, and memory.

An alternative logic that allows agents to hold inconsistent knowledge without knowing everything is the *bilattice logic* of Arieli and Avron [5], which itself is based on *Belnap's four-valued* logic [15, 16]. In bilattice logic, propositions can have, besides the truth values *true* and *false*, two more values, namely: (i) *both true and false* for handling inconsistency- this corresponds to the situation where several sources assign a different truth value to a sentence, and (ii) *neither true nor false*, for lack of information. It is interesting to note that the value *both true and false* in Belnap's logic can be traced back to Łukasiewicz's proposal [113] to address the *future contingent problem* in Aristotelian logic, which says that the truth of future events cannot be determined in the present unless we have complete information about future [51]. By adding modalities to Arieli and Avron's logic one can reason about different notions such as knowledge, belief, and time. Jung and Rivieccio [98] developed such an expansion of the bilattice logic of [5].

The approaches we mentioned so far only describe states of information and do not model *information change*. In the next section, we will discuss dynamics of knowledge.

1.1.3 Dynamics of information

An agent's knowledge may change in response to *informative events*. The following examples illustrate that *actions* may lead to knowledge change.

Example 1.1.8 Reconsider Example 1.1.1 and assume Anne tricks Bill and peeps at the card that is on the table, without Bill noticing it. In that case, Anne comes to know what the distribution of cards is, while Bill does not know that Anne knows the cards. In another scenario, suppose Anne walks towards the table, picks up the facedown card and looks at it without showing it to Bill, but he notices that Anne is looking at the card on the table. Similar to the previous scenario, Anne comes to know what the actual deal of the cards is, and Bill still does not know the distribution of the cards. The difference between these two scenarios lies in what Bill knows about Anne's knowledge about the cards. In the first scenario, Bill does not know that Anne knows the deal of the cards, while in the second scenario, Bill knows that Anne knows the deal of the cards, although he remains uncertain about the distribution of the cards himself. ⊣

Example 1.1.9 Reconsider Example 1.1.5, and assume Mona, who always speaks the truth, tells Leila "I prefer Wonder Woman over the other two movies, and you do not know that." As a result, Leila now knows that Wonder Woman is their favourite movie, and she is no longer ignorant about this fact. Consequently, she no longer has an incentive to manipulate the vote as her favourite movie will be selected if she votes sincerely. ⊣

There are two points in the above example that we would like to highlight. Firstly, Leila's state of information has changed due to an *announcement* that has been made by Mona. This announcement *reduces* Leila's uncertainty and accordingly has impact on her desire to manipulate. Announcements [135] are a relatively simple kind of actions that provide agents with new information that may resolve an agent's ignorance about a situation. Secondly, Mona has announced something that becomes false after the announcement. Before the announcement, Leila did not know Mona's preferences, but now that she knows, the statement is not true anymore. Motivated by these points, one of the objectives of the dissertation is to study knowledge-theoretic aspects of voting and manipulation, with *dynamic epistemic logic*.

Dynamic epistemic logic (DEL) [49, 11] is a family of logics for reasoning about change of knowledge. They expand the modal logic of knowledge with dynamic operators. One of the questions that is of interest in DEL is: how knowledge can be made known to an agent? This question falls under the topic *"quantifying over information change"* [45, 46, 86, 84, 31]. One logic that studies quantifying over information change is called *refinement modal logic*. Refinement modal logic (RML), introduced by Bozzelli et al. [31], is an extension of modal logic with operators that quantify over all the *refinements* of a model. A refinement intuitively corresponds to the result of an action that may change the agents' states of knowledge.

DEL is a well suited framework to describe the dynamics of knowledge, although as mentioned in Section 1.1.1, a downside to this approach is that it does not account for situations where agents have to cope with inconsistent knowledge (Examples 1.1.6, 1.1.7) or incomplete information (1.1.4). These situations are better analysed with *non-classical* logics such as *four-valued* logic. It raises the following research question: (i) how to formalise information change and quantifying over information change in the presence of incomplete and inconsistent information? Our methodology to answer this question is the duality theory in particular we apply the methods of [114, 105, 138] to refinement modal logic and the four-valued modal logic developed in [97].

1.2 Expressivity and bisimulation

The topic of Chapters 5 and 6 of this dissertation is finding an adequate notion of *bisimulation* for *weakly expressive* modal logics.

1.2.1 Kripke bisimulation and expressivity of basic modal logic

Modal languages are used to express properties of models. Different modal languages that are interpreted over the same class of models may differ in their expressive power. Expressivity is a measure of the power that a modal language has in distinguishing between different models. Let us provide an example which is borrowed from [26] with slight changes. Consider the models shown below.

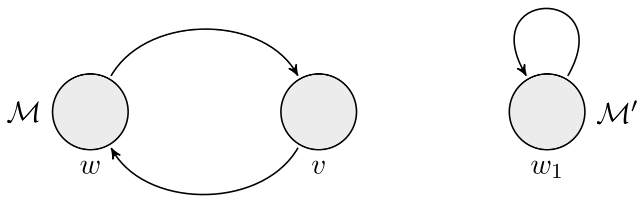

The model $\mathcal{M}'$ is reflexive, while the model $\mathcal{M}$ is not reflexive, and as reflexivity can be expressed as a first-order formula [25], it means that the models $\mathcal{M}$ and $\mathcal{M}'$ are distinguished by a first-order formula. However, $\mathcal{M}$ and $\mathcal{M}'$ satisfy the same formulas in the basic modal language, i.e., they are *modally equivalent*. Concerning this example, one may ask why these models are modally indistinguishable? Or more generally, one may ask when are two distinct models modally equivalent. A standard tool used to answer these questions is bisimulations [17]. It is well known that modal formulas are invariant under bisimulations, which means that if two states in a Kripke model are bisimilar (i.e., related by a bisimulation), then they are modally equivalent. The converse fails in general (see, e.g., [25, Example 2.23]), but it holds over the class of *image-finite Kripke models*. This result is referred to as the Hennessy-Milner Theorem for basic modal logic [93].

Another important link between bisimulation and expressivity of modal logic was established by Van Benthem's characterisation theorem [17]. This theorem says that every bisimulation invariant property of Kripke models that can be defined in first-order logic, is also definable in basic modal logic.

1.2.2 Weakly expressive modal logics

Modal languages differ in their power to express properties of the models they are interpreted in. We say that a modal language is weakly expressive, if everything that we can say in this language can be expressed in basic modal logic, but not vice versa. An example of such a modal language is given by contingency logic [62] in which modalities denote (non-)contingency of propositions.. A proposition is non-contingent if it is necessarily true or necessarily false, otherwise it is contingent. In an epistemic setting, "the proposition is non-contingent" becomes "the agent knows whether the proposition is true" and "the proposition is contingent" becomes "the agent is uncertain whether the proposition is true". The notion of "knowing whether" can be expressed in terms of "knowing that": an agent knows whether a proposition is true, if she knows that it is true or if she knows that it is false. Also, "knowing that" can be expressed as "knowing whether": an agent knows that a proposition is true, if the proposition is true and she knows whether the proposition is true. However, in the absence of the property of knowledge that known propositions are true, the notion of necessity (knowing that) can not necessarily be defined in terms of the notion of non-contingency (knowing whether). Contingency logic is thus less expressive than basic modal logic. Two different semantics are proposed for contingency logic, one based on Kripke models [62, 63], and the other one based on *neighbourhood* models [61]. The latter is motivated by the fact that non-contingency as a modal operator is not monotonic, and therefore contingency logic is not a normal modal logic. Neighbourhood models [37, 121] are a generalisation of Kripke models, and they have become the standard semantic tool for reasoning about non-normal modal logics.

Problem statement As mentioned above, the Hennessy-Milner theorem and Van Benthem's characterisation theorem show that Kripke bisimilarity matches the expressive power of the basic modal language over Kripke models. These results have been generalised to neighbourhood models and neighbourhood bisimilarity in [89]. However, for weakly expressive logics, such as contingency logic, the notions of Kripke bisimilarity and neighbourhood bisimilarity are too strong, meaning that two states that satisfy the same formulas may fail to be Kripke/neighbourhood bisimilar. To remedy this situation for contingency logic, Fan et al. [62] propose a notion of contingency bisimulation over Kripke models, and they show a Hennessy-Milner theorem and a characterisation theorem. However, their definition is found in an ad hoc manner, and it is not clear how to generalise contingency bisimulations to neighbourhood models. The following questions therefore arise: What is the right notion of neighbourhood contingency bisimulation? And more generally, given a weakly expressive logic, can we systematically define a matching notion of bisimulation?

Methodology and existing work To approach the above questions, we move to a higher level of abstraction and use the framework of coalgebra and coalgebraic modal logic. *Coalgebras* [141] can be viewed as an abstraction of state-based systems such as labelled transitions systems, Kripke frames, and neighbourhood structures [82]. Informally, a coalgebra is a set of worlds together with a transition map the type of which is parametric in the choice of functor. The theory of coalgebras is closely related to modal logic in two ways: Firstly, models of modal logic such as Kripke frames and neighbourhood structures can be represented as coalgebras. Hence, coalgebras generalise the traditional models of modal logic. Secondly, it is shown that modal logic is an adequate tool for reasoning about coalgebras, in the same sense that equational logic is the basic logic of algebras [104]. Indeed, researchers introduced a general theory known as *coalgebraic modal logic* [124, 131, 132] as a general framework in which modal logics for different types of structures can be developed in a uniform setting. The coalgebraic modal logic developed by Pattinson in [131] uses so-called predicate liftings to define a modal language. Informally, a predicate lifting can be seen as a generalisation of the necessity (possibility) modalities of basic modal logic. Coalgebraic modal logic comes with general methods and tools to prove soundness and completeness [131, 145], decidability [144], expressivity [132, 143], and develop correspondence theory [111].

Our motivation to use a coalgebraic framework is twofold. First, due to the above-mentioned parametricity in type, coalgebra allows for definitions and results to be developed uniformly for different types of structures. These results can then be instantiated for concrete classes of structures and their modal logics. Second, the theory of coalgebras comes with general notions of equivalence between states, namely, coalgebraic bisimulations [1, 141] and behavioural equivalence [104], as well as general results on bisimulation invariance and expressivity [132, 143].

Coalgebraic modal logic, as considered in [131], is invariant under behaviourally equivalence, however the converse does not hold in general. Pattinson in [132] proposed a condition on predicate liftings under which coalgebraic modal logic is *expressive*, meaning that if two states satisfy the same formulas, then they are behaviourally equivalent. Existing work in coalgebraic modal logic has focused on identifying conditions that ensure a modal language is expressive. The aim in this thesis is to turn things around and start from a weakly expressive coalgebraic modal logic, and find a notion of bisimulation that matches its expressivity. Our results therefore also contribute to the general theory of coalgebraic modal logic.

1.3 Organisation and contributions

We conclude the introduction by giving an overview of the chapters of this dissertation, their research themes, and a brief statement of their contributions.

Chapter 2 provides a novel framework for discussing manipulation under incomplete knowledge. Our main contributions in this chapter are: (i) the introduction of a modal logic of knowledge to account for a voter's uncertainty about other voters' preferences; (ii) addressing the question of how a reduction in uncertainty may affect manipulation. The main merit of these contributions is that they enable us to model *higher-order knowledge*. For example, we can model that voter 1 is uncertain about voter 2's preferences, and voter 1 knows that voter 3 knows voter 2's preferences. This makes it possible to model scenarios wherein all voters have the same preferences, but may have different uncertainty about the preferences of other voters.

Chapter 3 introduces an algebraic semantics for refinement action modal logic by following the methods presented in [114, 105]. Our main contribution is that we show that refinement action modal logic is sound and complete with respect to this algebraic semantics. This work is of interest as it is a first step towards developing non-classical counterparts of refinement modal logic.

Chapter 4 presents a dynamic epistemic logic for reasoning about information change in the presence of inconsistent and incomplete information. The logic we introduce in this chapter is an extension of the four-valued modal logic of [98] with dynamic modalities. The main contribution of this chapter is that we propose a sound and complete axiomatisation. This work bridges the frameworks of dynamic epistemic logic and many-valued modal logics [66, 98, 28]. It paves the way to study the mathematical foundations of dynamics of knowledge in non-classical settings.

Chapter 5 introduces the notion of *neighbourhood contingency bisimulation* for contingency logic interpreted over neighbourhood models. Our main contributions in this chapter are: (i) a Hennessy-Milner theorem for neighbourhood contingency bisimulation; (ii) characterisation of contingency logic over neighbourhood models as the bisimulation invariant fragment of first-order logic and of basic modal logic; (iii) showing that contingency logic has the Craig interpolation property.

Chapter 6 studies the notion of bisimulation developed in the previous chapter in the framework of coalgebraic modal logic. We introduce a notion of Λ-bisimulation for weakly expressive coalgebraic modal logics. The main technical result of this chapter is that we prove a Hennessy-Milner theorem for Λ-bisimulations.

1.4 Origin of materials

- Chapter 2 is based on:
  Zeinab Bakhtiari, Hans van Ditmarsch, and Abdallah Saffidine. *How does uncertainty about other voters determine a strategic vote* (2017). Under review.

- Chapter 3 is based on:
  Zeinab Bakhtiari, Hans van Ditmarsch, and Sabine Frittella. *Algebraic Semantics of Refinement Modal Logic.* In Lev D. Beklemishev, S. Demri, and A. Maté (eds.), Proceedings of the 11th Conference on Advances in Modal Logic (AiML 2016), Budapest, Hungary, August 30 - September 2, 2016. pp. 38–57, College Publications, 2016.

- Chapter 4 is based on two papers, where the latter is an extended version of the former:

  - Zeinab Bakhtiari and Umberto Rivieccio. *Epistemic Updates on Bilattices.* In W. van der Hoek, W. H. Holliday, and W. Wang (eds.), Proceedings of the 5th International Workshop, Logic, Rationality and Interaction (LORI 2015), Taipei, Taiwan, October 28-31, 2015. Lecture Notes in Computer Science, volume 9394, pp. 426–428, Springer, 2015.
  - Zeinab Bakhtiari, Hans van Ditmarsch, and Umberto Rivieccio. *Bilattice Logic of Epistemic Action and Knowledge* (2017). Under review.

- Chapter 5 is based on:
  Zeinab Bakhtiari, Hans van Ditmarsch, and Helle Hvid Hansen. *Neighbourhood Contingency Bisimulation.* In S. Ghosh and S. Prasad (eds.), Proceedings of the 7th Indian Conference on Logic and Its Applications - 7th Indian Conference, (ICLA 2017), Kanpur, India, January 5-7, 2017. Lecture Notes in Computer Science, volume 10119, pp. 48–63, Springer, 2017.

- Chapter 6 is based on:
  Zeinab Bakhtiari and Helle Hvid Hansen. *Bisimulation for Weakly Expressive Coalgebraic Modal Logics.* In F. Bonchi and B. König (eds.), Proceedings of the 7th Conference on Algebra and Coalgebra in Computer Science (CALCO 2017), Ljubljana, Slovenia, 14-16 June 2017. Leibniz International Proceedings in Informatics (LIPIcs), to appear, 2017.

2

Strategic voting and the logic of knowledge

Contents

2.1 Introduction

A well-known fact in social choice theory is that strategic voting, also known as manipulation, becomes harder when voters know less about the preferences of other voters. Standard approaches to manipulation in social choice theory [76, 142] as well as in computational social choice [13] assume that the manipulating voter knows the true preferences of other voters. Some approaches [52, 12] assume that voters have a probabilistic prior belief on the outcome of the vote, which encompasses the case where each voter has a probability distribution over the set of profiles. Coalitional manipulation was extended in [41] to contexts where manipulators have incomplete knowledge about the non-manipulators' votes. In some iterated voting settings, voters have incomplete knowledge of the other votes *at the next iteration*; this is the view taken by [120] and [119]. Still, we think that the study of strategic voting under complex belief states deserves more attention, especially when voters are uncertain about the uncertainties of other voters, i.e., when we model higher-order beliefs of voters.

An extreme case of uncertainty is when a voter is completely ignorant about other voters' preferences. In that case, if a manipulation under incomplete knowledge is defined in a pessimistic way, i.e., if it is said to be successful if it succeeds for all possible votes of other voters, voting rules may well be non-manipulable. For the special case where all other voters are non-strategic this is shown for most common voting rules in [41].

In this chapter, we model how uncertainty about the preferences of other voters may determine a strategic vote, and how a reduction in this uncertainty may change a strategic vote. We restrict ourselves to the case where each voter is uncertain about the number of well-described possible preference profiles, including the actual profile.

We also investigate the dynamics of uncertainty. The uncertainty reduction may be due to receiving information on voting intentions in polls or to voters informing other voters of their preferences. For simplicity we assume that received information is correct, or rather, we only model the consequences of incorporating new information after the decision to consider the information reliable. Such informative actions can then be modelled as truthful public announcements [135]. Does this information affect your strategic vote? There is a clear relation here to *safe manipulation* [150], where the manipulating voter announces her vote to a (presumably large) set of voters sharing her preferences but is unsure of how many will follow her.

We briefly survey existing approaches to representing incomplete knowledge about the preference of a voter. In the theory of voting, the voters' preferences over alternatives is modelled as a linear order [32]. The literature on possible and necessary winners [158, 91] represents incomplete knowledge about preferences as a collection of partial strict orders (one for each voter), while [91] formalises it as a collection of probability distributions, or a collection of sets of linear orders (one for each voter). Whereas the latter is more expressive (some sets of linear orders do not correspond to the set of extensions of a partial order), the former is more succinct. Ours is a more expressive modeling than those of [158, 91] because an uncertain profile can be any set of profiles. A set consisting of the two profiles $\{a \succ_1 b \succ_1 c, a \succ_2 b \succ_2 c\}$ and $\{b \succ_1 a \succ_1 c, b \succ_2 a \succ_2 c\}$ expresses uncertainty (ignorance) about which candidate voters 1 and 2 rank first, but also knowledge (certainty) that voters 1 and 2 have identical preferences — which is not possible in [91], and *a fortiori* also not in works on the possible winner problem [100] and more generally on voting under incomplete knowledge (see [30] for a survey). Of course, this mode of representation is also the least succinct of all. However, succinctness and complexity issues play no role in this chapter, where we focus on modeling and expressivity. Our representation of incomplete knowledge is independent from assumptions on the probability distribution (and therefore compatible with any assumption of that kind).

The main novelty of our proposal is that there are scenarios that cannot be seen as uncertainty between a number of given profiles: it may be that a voter cannot distinguish between two situations with identical profiles, because in the first case yet another voter has some uncertainty about the profile, but in the other case not. Such higher-order uncertainty is a common assumption in multi-agent modal logics, but to our knowledge this has not yet been investigated in voting theory.

Our investigation is restricted in various ways: (i) we model uncertainty and manipulability of individuals but not of coalitions, (ii) we model knowledge but not belief, and, in the dynamics, truthful announcements but not lying, (iii) we model incomplete knowledge (uncertainty) but not other forms of incompleteness, and (iv) as already said, we have not investigated complexity and succinctness issues. The reason for the restrictions (i) and (ii) is expository and not theoretical, anything we propose can be presented in that more general setting as well. In fact there are many scenarios in which voters may have incorrect beliefs about others' preferences, or where information changing actions are intended to deceive. I may incorrectly believe that you prefer a over b, whereas you really prefer b over a. I may tell you that I prefer a over b, but I may be lying. Such scenarios can also be modelled in epistemic logic, with the same tools and techniques as presented in this chapter. Restriction (iii) is more fundamental. An example is when the number of voters or candidates may be unknown. This relates to unawareness [57].

A link between epistemic logic and voting has first been given, as far as we know, in [39]—they use knowledge graphs to indicate that a voter is uncertain about the preferences of another voter. A more recent approach, within the area known as social software, is [130] which discusses games with knowledge manipulator agents. The works [41], [120] and [119] walk a middle way namely where equivalence classes are called information sets, as in treatments of knowledge and uncertainty in economics, but where the uncertain voter does not take the uncertainty of other voters into account, the main novelty of our proposal. We will come back on these works in later sections.

A different line of work, which has lead to a substantial amount of research, dating back to [64], consists in analyzing strategic voting using game-theoretic models, where actions available to voters are their votes, and where voters are assumed to have full knowledge of the preferences of other voters. However, the multiplicity of Nash equilibria makes this analysis difficult as soon as there are more than a few strategic voters. The case of few voters leads to interesting findings, such as the recent work by [53] on two-manipulator voting games for the Borda and k-approval voting rules. In this work too, voters' preferences are common knowledge. Our work can also be seen as a generalization of these voting games: we define equilibria for games with several *risk-averse* strategic voters, i.e., voters who cast their vote to minimise their worst possible consequences. We provide scenarios in which given the same profile, the equilibria are different, because the voters have different knowledge about other voters.

Our setting shares some similarity with robust mechanism design [21], which generalizes classical mechanism design by weakening the common knowledge assumptions of the environment among the players and the planner. In [21] uncertainty is modelled with information partitions. The main technical difference is that in our setting, as in classical social choice theory, preferences are ordinal, whereas in (robust) mechanism design preferences are numerical payoffs, which allows for payments (which we don't). A recent proposal on uncertainty in voting is [119], whose goals are similar to ours (when do equilibria exist, assume risk aversion) but whose methods are statistical (there is no higher-order uncertainty).

Organisation of the chapter. In Section 2.2 we recall standard voting terminology. In Section 2.3 we introduce knowledge profiles as a model for voters' uncertainty about preferences. Section 2.4 provides the links between knowledge and manipulation, and studies manipulation under uncertainty from game theoretical and voting theoretical perspective. In Section 2.5, we investigate the notion of dominant manipulation. Section 2.6 introduces the notion of equilibria in the presence of uncertainty, and provides examples to illustrate how we can determine equilibra when voters have uncertainty about preferences . In Section 2.7 we model actions that change uncertainty about preferences, such as a voter revealing its voting preferences. We prove that some forms of manipulation and equilibrium are preserved under such updates and others are not. Section 2.8 formalises epistemic voting terminology in a simple logic. In Section 2.9 we collect some conclusions and indicate further directions.

2.2 Voting

Assume a finite set $Ag = \{1, \ldots, n\}$ of n *voters* (or *agents*), and a finite set $\mathcal{C} = \{a, b, c, \ldots\}$ of m *candidates* (or *alternatives*). Voter variables are i and j, and candidate variables are x and y (and $x_1, x_2, \ldots$). We denote by $O(\mathcal{C})$ the set of linear orders on $\mathcal{C}$. A linear order over a set A is a binary relation which is irreflexive, transitive, complete, and antisymmetric.

Definition 2.2.1 (Vote, profile, voting rule) *Elements of $O(\mathcal{C})$ are referred to as* votes *or* preferences. *A* (voting) profile *for Ag is an element $P = (P_1, \ldots, P_n)$ of the product $O(\mathcal{C})^n$ where P_i is the vote cast by voter i in P. A* (resolute) voting rule *is a function $F : O(\mathcal{C})^n \to \mathcal{C}$ from the set of all profiles for Ag to the set of candidates.*

We write $\succ_i$ to denote voter's i preferences over candidates, where $\succ_i \in O(\mathcal{C})$. For example, we write $a \succ_i b$, if voter i prefers candidate a to candidate b. A vote by voter i which is equal to her preferences is a *sincere* vote, otherwise it is called an *insincere* vote. We write $\succ'_i, \succ''_i$, etc. to denote votes by voter i. Instead of $\succ'_i$ we may explicitly write $x_1 \succ'_i \cdots \succ'_i x_n$, or depict $\succ'_i$ vertically in a table. Profile variables are denoted by P, P', P'', etc. A *sincere profile* is a (unique) profile that consists of sincere votes for each voter.

Given a profile $P \in O(\mathcal{C})^n$, we define $P_{-i} = (P_1, \ldots, P_{i-1}, P_{i+1}, \ldots, P_n) \in O(\mathcal{C})^{n-1}$ be a profile without P_i. Thus we can write $P = (P_{-i}, P_i)$. If $P \in O(\mathcal{C})^n$ and $\succ'_i \in O(\mathcal{C})$, then $P' = (P_{-i}, \succ'_i)$ is the profile in which P_i is substituted by $\succ'_i$ in P.

For every $a, b \in \mathcal{C}$, and $i \in Ag$ with preference relation $\succ_i$, we define the binary relation $\succeq_i$ as follows: $a \succeq_i b$ iff $a \succ_i b$ or $a = b$. We say that voter i is *indifferent* between candidates a and b, if $a \succeq_i b$ and $b \succeq_i a$. Since the true preferences are linear orders, and we assume linear orders are irreflexive, voters are assumed to not be indifferent between any candidates.

A voting rule F determines which candidate wins the election — $F(P)$ is the *winner*. In case there is more than one *tied co-winner* we assume an exogeneously specified *tie-breaking mechanism*, that is a linear order $\succ$ over candidates.

Voters cannot be assumed to vote according to their true preferences. The voter may cast an insincere vote that leads to a better outcome than the sincere vote. Such a vote called a *strategic vote* or *manipulation.*

Throughout the present chapter, we assume the voters are *rational*; that is, they vote in a manner to obtain a most preferred outcome. It is often assumed that the rationality of all voters is common knowledge. We assume the reader is familiar with common voting procedures such as *plurality* voting and *Borda voting* (also called Borda count). For more details on voting rules we refer to [30, 27].

Definition 2.2.2 (Successful manipulation) *Let $i \in Ag$, $P \in O(\mathcal{C})^n$ with $P_i = \succ_i$, and $\succ'_i \in O(\mathcal{C})$. If $F(P_{-i}, \succ'_i) \succ_i F(P)$, then $\succ'_i$ is a* successful manipulation *by voter i in profile P.* ⊣

The combination of a profile P and a voting rule F defines a strategic game: a player is a voter, an individual strategy for a player is a vote, a strategy profile (of players) is therefore a profile in our defined sense (of voters), and the preferences of a player among the outcomes is according to his (true) preferences: given profiles P', P'', voter i prefers outcome $F(P')$ over outcome $F(P'')$ in the game theoretical sense iff (in the voting sense) $F(P') \succ_i F(P'')$, where $\succ_i$ is i's (true) preferences. Let us note that throughout the present chapter we only consider pure strategies. The relevant notions of dominance and equilibrium are:

Definition 2.2.3 (Domination) *Let $\succ'_i$ and $\succ''_i$ be two votes for voter i. Then*

1. *vote $\succ'_i$* strongly dominates *$\succ''_i$, if for all $P_{-i} \in O(\mathcal{C})^{n-1}$: $F(P_{-i}, \succ'_i) \succ_i F(P_{-i}, \succ''_i)$.*
2. *vote $\succ'_i$* weakly dominates *$\succ''_i$, if for all $P_{-i} \in O(\mathcal{C})^{n-1}$: $F(P_{-i}, \succ'_i) \succeq_i F(P_{-i}, \succ''_i)$, and for some $P_{-i} \in O(\mathcal{C})^{n-1}$: $F(P_{-i}, \succ'_i) \succ_i F(P_{-i}, \succ''_i)$.* ⊣

Now we can define *dominant manipulation* for a voter as a vote that is a manipulation and dominates all other votes for voter i.

Definition 2.2.4 (Dominant manipulation) *Let $i \in Ag$.*

- *The vote $\succ'_i \in O(\mathcal{C})$ is a* weakly dominant manipulation *for voter i if*
  1. *$\succ'_i \neq \succ_i$, and*
  2. *for all $\succ''_i \in O(C)$ such that $\succ''_i \neq \succ'_i$: $\succ'_i$ weakly dominates $\succ''_i$.*
- *The vote $\succ'_i \in O(\mathcal{C})$ is a* strongly dominant manipulation *for voter i if*
  1. *$\succ'_i \neq \succ_i$, and*
  2. *for all $\succ''_i \in O(C)$, $\succ'_i$ strongly dominates $\succ''_i$.* ⊣

A weakly (strongly) dominant manipulation corresponds to the natural notion of weakly (strongly) dominant strategy in game theory [110]. (Dominant manipulation *with respect to a set of profiles* will be discussed in Section 2.5.)

Definition 2.2.5 (Equilibrium profile) *A profile P is an* equilibrium profile *iff no voter has a successful manipulation in P.* ⊣

An equilibrium profile corresponds to a Nash equilibrium (in game theory). An equivalent way of defining equilibrium profile is: A profile P is an equilibrium profile iff

$$\text{For all } i \in Ag, \text{ and for all } \succ'_i \in O(\mathcal{C}) : F(P) \succeq_i F(P_{-i}, \succ'_i). \tag{2.1}$$

We show that the equation (2.1) is equivalent to Definition 2.2.5. Consider the negation of (2.1), i.e., there is a voter i and some vote $\succ'_i \in O(\mathcal{C})$ such that $F(P) \prec_i F(P_{-i}, \succ'_i)$, i.e.,

$$F(P_{-i}, \succ'_i) \succ_i F(P). \tag{2.2}$$

Note that since $\succ_i$ is irreflexive, this implies that $\succ'_i \neq \succ_i$. Thus, (2.2) is equivalent to saying that the vote $\succ'_i$ is a successful manipulation for voter i in profile P. In other words, the condition (2.1) holds iff no voter has a successful manipulation in P.

2.3 Knowledge profiles

We model uncertainty in voting as incomplete knowledge about voters' preferences. The structures we use to represent uncertainty and the terminology we use to describe uncertainty that we introduce in this section, are fairly standard in modal logic [58, 49], but this application to social choice theory is novel. From a modal logic point of view, the novelty consists in taking models with *profiles* instead of *valuations of propositional variables*, whereas from a social choice theory point of view, the novelty consists in making a relational structure consisting of profiles a semantic primitive, instead of (only) a single profile or (only) a set of profiles. There are different ways in which a profile may correspond to a valuation. This is explained in Section 2.8. In order to allow for the definition of dominance and of equilibria under uncertainty, we further require that voters know their own preferences.

Definition 2.3.1 (Knowledge profile) *Let $O(\mathcal{C})^n$ be the set of all profiles for a set $Ag = \{1, \ldots, n\}$ of n voters. A* profile model *(Kripke model) is a structure $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$, where S is a* domain *of abstract objects called states, $\sim_i$ is an* indistinguishability relation *that is an equivalence relation for $i = 1, \ldots, n$, and $V : S \to O(\mathcal{C})^n$ is a* profile function *that assigns a profile to each state such that $s \sim_i t$ implies $V(s)_i = V(t)_i$, i.e., voters know their preferences. We let $[s]_{\sim_i}$ denote $\{t \mid s \sim_i t\}$, and we let $V([s]_{\sim_i})$ denote $\{V(t) \mid s \sim_i t\}$. A* knowledge profile *(pointed Kripke model) is a pointed structure $(\mathcal{M}, s)$ where $\mathcal{M}$ is a profile model and s is a state in the domain of $\mathcal{M}$.* ⊣

Let $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$ be a profile model such that $s, t \in S$. If $s \sim_i t$, and $V(s) = P$ and $V(t) = P'$, then we say voter i is *uncertain (ignorant)* if the profile is P or P'. Instead of saying that "voter i is uncertain if ..." we also say "voter i does not know that ...". Statements such as "P is the profile", "i prefers a over b", "voter i does not know that... ", "voter i knows j's preferences" and "voter j knows that i knows her preferences" are called *propositions about profiles*. In Section 2.8 we present the logic of knowledge and voting formally, with a logical language instead of natural language, in which we inductively introduce "*proposition about profiles*". However, in the present section by "propositions about profiles" we simply mean the kind of statements that we mentioned above.

The next definition shows how we can model knowledge and uncertainty in terms of knowledge profiles.

Definition 2.3.2 (Knowledge and ignorance) *Let $(\mathcal{M}, s)$ be a knowledge profile, where $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$. Let φ be a proposition about profiles. Voter i* knows *φ in $(\mathcal{M}, s)$, iff φ is true in all $t \in S$ such that $s \sim_i t$. Voter i* considers possible that *(or* does not know that not*) φ in $(\mathcal{M}, s)$, iff φ is true in some $t \in S$ such that $s \sim_i t$; if, in that case, there is an additional state $u \in S$ with $s \sim_i u$ in which φ is false, then we say that i* does not know whether *(or* is uncertain about*, or* is ignorant about*) φ.* ⊣

The following Section 2.3.1 will provide examples. It will also demonstrate why profiles are not sufficient for our knowledge representation and why we need states: *different states may be assigned the same profile, but have different knowledge properties.* In scenarios where different states are always assigned different profiles, we can say that the uncertainty of a voter is (only) about a collection of profiles. In scenarios where different states are assigned the same profile, the set $V([s]_{\sim_i})$ of profiles that voter i considers possible is potentially smaller than the set $[s]_{\sim_i}$ of states that i considers possible.

In order to explain the interaction between knowledge and preference, very simple settings with few voters and little uncertainty are sufficient. More realistic voting scenarios for larger populations of voters that also involve uncertainty, quickly become very complex because of that additional feature. Our running examples, as in the next section, will therefore be of the first simple kind.

Partial preferences can not be expressed in our framework. In particular, uncertainty between $a \succ_i b \succ_i c$ and $b \succ_i a \succ_i c$ does not mean that voter i is indifferent between candidates a and b. Uncertainty between $a \succ_i b \succ_i c$ and $b \succ_i a \succ_i c$ means that ($a \succ_i b \succ_i c$ or $b \succ_i a \succ_i c$) is true. This entails ($a \succ_i b$ or $b \succ_i a$), which is equivalent to (not ($b \succeq_i a$ and $a \succeq_i b$)). That is the opposite of indifference between a and b, as that means ($b \succeq_i a$ and $a \succeq_i b$).

2.3.1 Example

Consider two voters, Leela (1) and Sunil (2), who are the children of Devi, and who 'vote' for an animated movie to see tonight before bedtime; where the choice is between a (Alice in

Wonderland), b (Brave), and c (Cars). Let us be traditional: Alice in Wonderland is a girl's movie, Cars is a boy's movie, and Brave (about a bow-and-arrow wielding girl hero in the Scottish highlands) is both. Leela's preferences are $a \succ_1 b \succ_1 c$ and Sunil's preferences are $c \succ_2 b \succ_2 a$. (Mother Devi will feature only later as the tie-breaking mechanism — after all, they must go to bed at some stage.) Sunil (2), who does not care about traditional gender-based preferences, is uncertain if Alice in Wonderland is Leela's top or bottom preference, and if Cars is Leela's bottom or top preference. Leela's (1) uncertainty is of the more interesting kind: she knows what Sunil likes, but she is uncertain if Sunil is uncertain or not.

We model this as the following knowledge profile $(\mathcal{M}, t)$ consisting of three states s, t, u for two voters 1 and 2. Profile P is assigned to state s, in which $a \succ_1 b \succ_1 c$ and $c \succ_2 b \succ_2 a$, etc. States that are indistinguishable for a voter i are linked with an i-labelled edge. The partition for 1 on the domain is therefore $\{\{s,t\},\{u\}\}$, and the partition for 2 on the domain is $\{\{s\},\{t,u\}\}$.

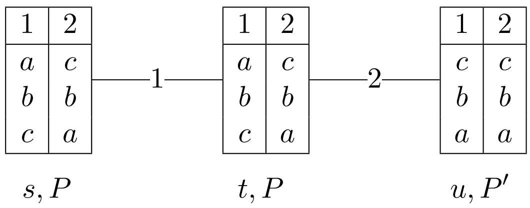

States s and t have been assigned the same profile P but have different epistemic properties. In s, voter 2 knows that 1 prefers a over c, whereas in t voter 2 does not know that. We list some relevant propositions that are true in the actual state t:

- *Leela prefers Alice in Wonderland over Cars.*
  This is true, because $a \succ_1 c$ in t.

- *Sunil does not know that Leela prefers Alice in Wonderland over Cars.*
  This is true, because $t \sim_2 u$, and $a \succ_1 c$ is false in u.

- *Leela knows what Sunil likes, but Leela is uncertain whether Sunil knows what Leela likes (i.e., she is uncertain if Sunil is uncertain [about Leela's preferences] or not).*
  The main observation here is that in u, Sunil knows that Leela's preferences are $c \succ_1 b \succ_1 a$, whereas in t, Sunil does not know that Leela's preferences are $c \succ_1 b \succ_1 a$ (because $t \sim_2 u$, and $a \succ_1 b \succ_1 c$ in u).

- *Leela and Sunil know their own preferences.*
  This can also be checked easily.

2.4 Manipulation and knowledge

In this section we present notions linking knowledge and successful manipulation. After that we will present game theoretical and voting theoretical notions for manipulation under such uncertainty.

In a knowledge profile it may be that a voter can successfully manipulate the vote but does not know that, because she considers it possible that another profile is the case, in which she cannot successfully manipulate the vote. Such scenarios call for more refined notions of manipulation that also involve knowledge. They can be borrowed from the knowledge and action literature [18, 96]. We now introduce these notions.

Definition 2.4.1 (Knowledge of successful manipulation) *Let $(\mathcal{M}, s)$ be a knowledge profile.*

1. *Voter i can successfully manipulate $(\mathcal{M}, s)$ if she can successfully manipulate the profile $V(s)$.*
2. *Voter i considers it possible that she can successfully manipulate $(\mathcal{M}, s)$ if there is a t such that $s \sim_i t$ and she can successfully manipulate $V(t)$.*
3. *Voter i knows 'de dicto' that she can successfully manipulate $(\mathcal{M}, s)$, if for all t such that $s \sim_i t$ she can successfully manipulate $V(t)$.*
4. *Voter i knows 'de re' that she can successfully manipulate $(\mathcal{M}, s)$ if there is a vote $\succ_i'$ such that for all t such that $s \sim_i t$, $\succ_i'$ is a successful manipulation for profile $V(t)$.* $\dashv$

Informally, the above definition says that:

1. Given a knowledge profile $(\mathcal{M}, s)$ where $V(s) = P$, if voter i can successfully manipulate P, then voter i can also successfully manipulate $(\mathcal{M}, s)$. This is because manipulation is defined with respect to the profile of the actual state of the knowledge profile. It is not defined with respect to voters' uncertainty about the profile. Although voters may be uncertain about what the profile is, this does not affect that P is the actual profile.
2. In our modelling, if the voter can manipulate P, she always considers it possible that she can manipulate P. This is a consequence of modelling uncertain knowledge instead of uncertain belief. We can easily imagine scenarios in which she considers it possible that she can manipulate, but where in fact she cannot manipulate. This can happen if she cannot successfully manipulate the actual profile, but there is a state that she considers possible in which she can successfully manipulate. As already said, for expository reasons we only model knowledge and not belief.
3. De dicto knowledge manipulation refers to a curious situation where in all states that the voter considers possible there is a successful manipulation, but where, unfortunately, this is not the same strategic vote in all such states. So she knows that she has a successful manipulation, but she does not know what the manipulation is. In other words, in 'de dicto' manipulations the voter does not seem to have the *ability* to manipulate the election. It is akin to 'game of chicken' type equilibria in game theory [128]. Therein, for each strategy of a player there is a complementary strategy of the other player such that the pair is an equilibrium. The existence of such equilibria cannot be guaranteed without coordination. 'De dicto' manipulations therefore do not seem to be of practical interest. Instead of a problem, we consider this a feature: the interest of voters in such a case is to obtain additional information that reduces the uncertainty (such as informing each other about their preferences), in order to get to know that successful manipulation is possible. Such uncertainty reduction is modelled in Section 2.7. Section 2.4.1 illustrates 'de dicto' manipulability.
4. De re knowledge of successful manipulation is a stronger form of knowing. If voter i knows 'de re' that she can manipulate the election with vote $\succ_i'$, she has the ability to manipulate, namely by strategically voting $\succ_i'$ in all states that she considers possible. The reader informed in voting theory will recognize this as a strongly dominant manipulation with respect to the set of profiles $V([s]_{\sim_i})$. As we do not wish to mix up epistemic and game

theoretical considerations, we will discuss this matter (and related matters) in a separate Section 2.5.

5. We assume that voters know their own preferences in profile models. The successful manipulations in all states she considers possible are all with respect to the *same* true preferences. To define knowledge of successful manipulation this is irrelevant(a voter can have the same successful manipulation but with respect to two different but indistinguishable conflicting votes), but for the relation with game theoretical notions of dominance and equilibrium, this is crucial.

Consider the profile model consisting of the domain $S = O(\mathcal{C})^n$ of all profiles, so we can identify states s with their profiles $P = V(s)$, and such that, whatever the profile is, all voters only know their own preferences, so that the accessibility relation is defined as $P \sim_i P'$ iff $P_i = P'_i$. In this (unique) model it is *common knowledge* that voters only know their own preferences, and as voters know their preferences, i.e., their *local state*, this is the interpreted system [58] consisting of all *global states* (i.e., profiles) known as a *hypercube* [112].

Proposition 2.4.2 *Given the 'hypercube' profile model in which voters only know their own preferences. Then 'de re' and 'de dicto' knowledge of successful manipulation is impossible for plurality voting.* ⊣

Proof We start with the 'de re' case. The 'de dicto' case is an inessential variation of the 'de re' case. Let P be the profile. We may assume that there are at least two candidates and at least two voters, as otherwise successful manipulation is not possible anyway. We first consider the case that there are at least three voters. Assume towards a contradiction that voter i (with sincere vote $\succ_i = P_i$) knows 'de re' that $\succ'_i$ is a successful manipulation for her. This means that for all P' such that $P'_i = \succ_i$, $F(P'_{-i}, \succ'_i) \succ_i F(P')$. This includes the profile P'' in which all other voters j have the same preferences as i's sincere preferences, i.e., for all $j \in Ag$: $P''_j = \succ_i$. In the profile P'', i's preferred candidate would have won by (unanimous) majority vote, contradicting the assumption that $\succ'_i$ is a successful manipulation since this implies that the candidate $F(P''_{-i}, \succ'_i)$ is strictly more preferred by i than $F(P'')$. It does not matter what the tie-breaking preferences are, as a majority of two voters already overrules the tie. The other boundary case is when there are exactly two voters, such that voting for the winner of the tie-breaking preference is dominant. But although this is a manipulation, if the voter's preferred candidate is different from the tie's preferred candidate, it is *not* a successful manipulation. In the 'de dicto' case there is a successful manipulation for each profile that voter i considers possible (where the manipulation may depend on the profile). As this also applies to the profile where all other voters j have the same preferences as i's sincere preferences, we again derive a contradiction. □

This result can be generalized to many other voting rules, such as the Borda voting rule in the next example section. However, under conditions of anonymity (i.e., when the voting rule is a symmetric function), it does not hold for all voting rules. Consider two voters $1, 2$ and two candidates a, b and and (a rather undemocratic) voting rule F defined as $F(\{a \succ_1 b, a \succ_2 b\}) = b$ and otherwise a, i.e., $F(\{a \succ_1 b, b \succ_2 a\}) = F(\{b \succ_1 a, a \succ_2 b\}) = F(\{b \succ_1 a, b \succ_2 a\}) = a$. Observe that the voting function is symmetric. Let the true preferences of voter 1 be $a \succ_1 b$ and of voter 2 be $a \succ_2 b$ and assume that the voters only know their own preferences. Then voter 1 knows that $b \succ_1 a$ is successful manipulation and voter 2 knows that $b \succ_2 a$ is a successful manipulation.

2.4.1 De dicto knowledge with Borda voting

We consider manipulation in the Borda voting rule. In Borda voting, the points for each candidate in each vote are added, and the candidate with the highest sum wins, modulo the tie-breaking preference. The preferred candidate gets 3 points, the 2nd choice 2 points, etc. Consider three agents, four candidates, and two profiles P and P' that are indistinguishable for agent 1, but that agents 2 and 3 can tell apart; as follows.

1	2	3
c	d	b
b	a	d
a	c	c
d	b	a

P

—1—

1	2	3
c	d	b
b	a	a
a	c	c
d	b	d

P'

The true preferences of the voters is given in profile P. There is also a tie-breaking preference $b \succ c \succ d \succ a$. The difference between the profiles P and P' is that 3 prefers d over a in P but a over d in P'. We prove the following:

> Voter 1 can successfully manipulate the election if the profile is P, and voter 1 can successfully manipulate the election if the profile is P', but the manipulation for P gives her a worse outcome for P', and the manipulation for P' gives her a worse outcome for P.

Therefore voter 1 is not effectively able to successfully manipulate the outcome of the election. She only knows 'de dicto' that she can successfully manipulate this election. She does not know how to vote in order to achieve that.

First, the outcome when all three voters give their true preferences. We write (x, y, z, w) when there are x points for a, y for b, z for c, w for d.

profile	count	observation	outcome
P	$(3, 5, 5, 5)$	b, c, d are tied	b
P'	$(5, 5, 5, 3)$	a, b, c are tied	b

Voter 1 can manipulate P or P' by downgrading b. But this is tricky, because it comes at the price of making a or d, or both, get a higher count than her preferred alternative c. This price is indeed too high:

In P, voter 1 can achieve a better outcome by the vote $\succ'_1$ defined as $c \succ'_1 a \succ'_1 b \succ'_1 d$. Let $Q = (P_{-1}, \succ'_1)$, and $Q' = (P'_{-1}, \succ'_1)$.

profile	count	observation	outcome
Q	$(4, 4, 5, 5)$	c, d are tied	c
Q'	$(6, 4, 5, 3)$		a

Although voter 1 prefers the winner in Q over the winner in P, the winner in Q' is less preferred by her than the winner in P'.

In P', voter 1 can achieve a better outcome by $\succ''_1$ defined as $c \succ''_1 d \succ''_1 b \succ''_1 a$. Let $R = (P'_{-1}, \succ''_1)$, and $R' = (P_{-1}, \succ''_i)$.

profile	count	observation	outcome
R	$(2, 4, 5, 7)$	1's least preference	d
R'	$(4, 4, 5, 5)$	c, d are tied	c

Now, voter 1 prefers the winner in R' over the winner in P', but the winner in R is less preferred by her than the winner in P.

For the record, these are the winners for all different votes for voter 1 where c is most preferred (we recall that $1 : cbad$ means $c \succ'_1 b \succ'_1 a \succ'_1 d$, etc.).

profile	$1 : cbad$	$1 : cabd$	$1 : cdba$	$1 : cadb$	$1 : cdab$	$1 : cbda$
P	$b(3,5,5,5)$	$c(4,4,5,5)$	$d(2,4,5,7)$	$d(4,3,5,6)$	$d(3,3,5,7)$	$d(2,5,5,6)$
P'	$b(5,5,5,3)$	$a(6,4,5,3)$	$c(4,4,5,5)$	$a(6,3,5,4)$	$c(5,3,5,5)$	$b(4,5,5,4)$

2.5 Dominance and knowledge

The notions of successful manipulation and dominant manipulation are unrelated to uncertainty over the profile. To model such uncertainty, the notion called dominant manipulation *with respect to an information set* (i.e., with respect to an $\sim_i$ equivalence class for voter i) is used in voting theory. Let us quote the clear description in [41]:

> "We suppose the knowledge of the manipulator is described by an information set E. This is some subset of possible profiles of the non-manipulators which is known to contain the true profile. Given an information set and a pair of votes U and V, if for every profile in E, the manipulator is not worse off voting U than voting V, and there exists a profile in E such that the manipulator is strictly better off voting U, then we say that U dominates V." [41, page 1]

In our terminology this becomes the following that matches the prior definitions of successful and (weakly) dominant manipulation in Section 2.2.

Definition 2.5.1 (Domination w.r.t. a set of profiles) *Let a voter $i \in Ag$ be given with true preference $\succ_i \in O(\mathcal{C})$, and let $\mathfrak{P}$ be a collection of profiles $P_{-i} \in O(\mathcal{C})^{n-1}$ that contains the profile of sincere votes for all players different from i.*

- *vote $\succ'_i \neq \succ_i$ strongly dominates $\succ''_i$ with respect to $\mathfrak{P}$, if for all $P_{-i} \in \mathfrak{P}$: $F(P_{-i}, \succ'_i) \succ_i F(P_{-i}, \succ''_i)$.*
- *vote $\succ'_i \neq \succ_i$ weakly dominates $\succ''_i$ with respect to $\mathfrak{P}$, if for all $P_{-i} \in \mathfrak{P}$: $F(P_{-i}, \succ'_i) \succeq_i F(P_{-i}, \succ''_i)$, and for some $P_{-i} \in \mathfrak{P}$: $F(P_{-i}, \succ'_i) \succ_i F(P_{-i}, \succ''_i)$.* ⊣

Similar to Def. 2.2.4, we can define dominant manipulation with respect to a set of profiles as follows:

Definition 2.5.2 (Dominant manipulation w.r.t a set of profiles) *Let a voter $i \in Ag$ be given with true preference $\succ_i \in O(\mathcal{C})$, and let $\mathfrak{P}$ be a collection of profiles $P_{-i} \in O(\mathcal{C})^{n-1}$ that contains the profile of sincere votes for all players different from i. Vote $\succ'_i$ is a* (weakly) dominant manipulation *for voter i with respect to $\mathfrak{P}$, if for all $\succ''_i \in O(\mathcal{C})$ such that $\succ''_i \neq \succ'_i$, vote $\succ'_i$ weakly dominates $\succ''_i$ with respect to $\mathfrak{P}$.* ⊣

For strong dominance we require that for all $\succ''_i \neq \succ'_i$, vote $\succ'_i$ strongly dominates $\succ''_i$ with respect to $\mathfrak{P}$. We can immediately observe that (i) a dominant manipulation for voter i with true preference $\succ_i$ in the set $O(\mathcal{C})^n$ of all profiles P corresponds to a dominant strategy in the standard game theoretical sense, i.e., to the dominant manipulation of Def. 2.2.4; and that (ii) a dominant manipulation in the singleton collection $\mathfrak{P} = \{P\}$ of true preferences for all voters (so including $\succ_i$) is a successful manipulation by voter i in profile P, as in Def. 2.2.2.

Observation (i) seems to suggest that dominant votes are about uncertainty. This is in a sense true, but only to the extent that strategic uncertainty (what action will my opponent choose given his freedom to act) is independent from epistemic uncertainty (what action will my opponent choose given my uncertainty about his nature).

Definition 2.2.4 of dominant manipulation (in the game theoretical sense) and the definition of dominant manipulation with respect to the set of all profiles containing the manipulator's sincere vote (in the sense of [41]) coincide in that set of profiles. Also, they coincide in the fact that the preferences of other voters play no role in determining whether the property holds. However, for notions other than dominance, such as equilibrium, in which the voter may reason about the preferences of other voters, the above difference matters a great deal; in a way, it is all that matters. That will be presented in the next section.

We conclude this subsection with some relations between dominance and knowledge. Given a knowledge profile $(\mathcal{M}, s)$, we can define a notion of knowledge relative to the profiles $V([s]_{\sim_i})$ indistinguishable for a voter, that captures the notion of dominant manipulation for information set $V([s]_{\sim_i})$ in [41]:

Definition 2.5.3 (Weak knowledge of successful manipulation) *Let $(\mathcal{M}, s)$ be a knowledge profile. Voter i with true preference $\succ_i$ weakly knows that she can successfully manipulate $(\mathcal{M}, s)$ if there is a vote $\succ_i'$ such that for all t such that $s \sim_i t$, $F(V(t)_{-i}, \succ_i') \succeq_i F(V(t))$, and for some t such that $s \sim_i t$, $\succ_i'$ is a successful manipulation for profile $V(t)$, i.e., $F(V(t)_{-i}, \succ_i'$ $) \succ_i F(V(t))$.* ⊣

This is a 'de re' notion of knowledge, although slightly different from the one in Def. 2.4.1:

Proposition 2.5.4 *If voter i knows 'de re' that she has a successful manipulation, then this is equivalent to her having a* strongly *dominant manipulation in the set $V([s]_{\sim_i})$.* ⊣

Some (but not all) of our results for knowledge of manipulation also hold for weak knowledge of manipulation. Our proof of Proposition 2.4.2 stating that knowledge of successful manipulation is impossible in the hypercube profile model (maximal ignorance of other voters' preferences) cannot be proved in the same way for weak knowledge of successful manipulation. For that, see [41], where however their setting is different, as they do not consider the hypercube but the model in which one designated voter, the manipulator, is ignorant of the preferences of all other voters, the non-manipulators (that are assumed to know the profile).

Knowledge has the properties of positive and negative introspection: if a voter i knows a proposition, then she also knows that she knows the proposition, and if a voter i does not know a proposition, then she also knows that she does not know it. So a voter i who has (weak) knowledge of successful manipulation also knows that she has (weak) knowledge of successful manipulation.

In a multi-agent setting in which there are several equivalence classes for voter i, each corresponding to a vote that i can cast, we can distinguish a profile model in which a voter has knowledge of successful manipulation in all her equivalence classes from a model in which she has knowledge of successful manipulation only in the actual equivalence class (but maybe not in all of her equivalence classes). We do not know if voting theorists have considered a similar distinction between *local dominant manipulation* (the actual information set) and *global dominant manipulation* (all information sets for that voter).

2.6 Equilibrium and knowledge

Determining equilibria under incomplete knowledge comes down to decision making under incomplete knowledge. Therefore we have to choose a decision criterion. Expected utility makes no sense here, because we didn't start with probabilities over profiles in the first place, nor with utilities. In the absence of prior probabilities, the following three criteria make sense. (*i*) The *insufficient reason* (or *Laplace*) criterion is a probability distribution over states that assigns equal probability to all possible states in a given equivalence class. This criterion was used in [4] to determine equilibria of certain (Bayesian) games of imperfect information. (*ii*) The *minimax regret* criterion selects the decision minimizing the maximum utility loss, taken over all possible states, compared to the best decision, had the voter known the true state. (*iii*) The *pessimistic* (or *Wald*, or *maximin*) criterion is a decision making rule that compares decisions according to their worst possible consequences. This rule states that the decision maker should take the strategy whose worst outcome is better than the worst outcome of the other strategies. The latter criterion, which we also call *risk aversion*, fits well our probability-free and utility-free model; this was also the criterion chosen in [41, 120, 119] ([119] also considers the minimax-regret criterion). The only assumption here in this chapter is that the probability distribution is positive in all states. In the rest of the chapter we work with criterion (iii). (Pessimistic, optimistic, and yet other criteria only assuming positive probability are applied to social choice settings in [130]. We think their interesting results can be modelled as games using our setting.)

In the presence of knowledge, and the assumption that voters know their own preferences (so that, in game theoretical terms, a voter's preferences are uniform throughout any of her equivalence classes, the definition of an equilibrium extends naturally. For each agent, the combination of an agent i and an equivalence class $[s]_{\sim_i}$ for that agent (for some state s in the knowledge profile) defines a so-called virtual agent: we model these imperfect information games as Bayesian games [90]. Thus, agent i is multiplied in as many virtual agents as there are equivalences classes for $\sim_i$ in the model. Each virtual agent has the same set of strategies as the 'original' agent.

In our setting we can almost think of these equivalence classes as sets of indistinguishable profiles. Almost but not quite: we recall that two states in the same equivalence class may be assigned the same profile but have different knowledge properties.

An equilibrium is then a profile of strategies such that none of the virtual agents has an interest to deviate. An alternative way to present a Bayesian game that does not use virtual agents, applied in [4], is to stick to the agents one already has but to change the set of strategies. Instead of each voter ('virtual agent') choosing a vote ('strategy') among the set of votes, in each equivalence class (information set), we have each voter choosing a *conditional vote* among the larger set of conditional votes, where the conditions are the equivalence classes for the agents. This presentation we will now follow in the definition below. For risk-averse voters knowing their own preferences we can effectively determine if a conditional profile is an equilibrium without taking probability distributions into account, unlike in the more general setting of Bayesian games that it originates from. (The difference with [4] is that they use the 'insufficient reason' decision criterion: uniform random choice between possible states.)

Definition 2.6.1 (Pessimistic successful manipulation) *Let $i \in Ag$ and let $\mathfrak{P} = \{P \in O(\mathcal{C})^n : P_i = \succ_i\}$ be the set of all profiles in which voter i votes sincerely. Assume $P \in \mathfrak{P}$ gives the worst (pessimistic) outcome in $\mathfrak{P}$ for voter i. We say that (risk-averse) voter i has a* pessimistic successful manipulation $\succ'_i$ in $\mathfrak{P}$ *if $\succ'_i$ is a successful manipulation in P and gives no worse outcome for any other profile in $\mathfrak{P}$, i.e.: for all $P' \in \mathfrak{P}$, $F(P'_{-i}, \succ'_i) \succeq_i F(P_{-i}, \succ'_i)$. A* *pessimistic successful manipulation for voter i in an equivalence class $[s]_{\sim_i}$ of a knowledge profile $(\mathcal{M}, s)$ is a pessimistic successful manipulation with respect to the set $V([s]_{\sim_i})$ of profiles.* ⊣

Definition 2.6.2 (Conditional vote, conditional profile, conditional equilibrium) *Let $\mathcal{M} = (S, \{\sim_1, \dots, \sim_n\}, V)$ be a profile model. For each voter i, a* conditional vote *is a function $[\succ]_i : S/{\sim_i} \to O(\mathcal{C})$ that assigns a vote to each equivalence class for that voter. A* conditional profile *is a collection of n conditional votes, one for each voter. A* conditional voting game *is then a (standard) strategic game in which voters declare conditional votes.* ⊣

In order to set the stage for examples of the next section, we explain how we determine the payoffs of the voters and equilibria under uncertainty. In the situation without uncertainty, given n voters, a profile and a voting rule determine a winner. In the presence of uncertainty, the outcome of a conditional profile consisting of conditional votes is an n-tuple of tuples or vectors $(x_i^1, \dots, x_i^m)$, where voter i has m equivalence classes and where $x_i^1, x_i^2, ..., x_i^m$ are the payoffs associated with each of the equivalence classes. For example, let $\mathcal{C} = \{a, b\}$, and let a voter i have two equivalence classes x and y. If she prefers a over b in x and b over a in y, then the vectors $(0, 1)$ means that in x, voter i gets payoff 0, and in y she gets payoff 1, and the vector $(1, 0)$ means that in x, she gets payoff 1 and in y she gets 0. The vectors $(x_i^1, \dots, x_i^m)$ are unordered, so we have to compute equilibria differently. For example, we cannot say which of $(0, 1)$ and $(1, 0)$ voter i prefers. But we can say that virtual voter (i, x) prefers the second (in which she gets 1) over the first, and that virtual voter (i, y) prefers the first over the second. This is the Bayesian game computation of equilibrium, where we determine manipulability for each virtual agent. Therefore, in the definition we did not write 'A conditional profile is an equilibrium iff no agent has a successful manipulation,' but '(...) if no agent has a pessimistic successful manipulation *in any of its equivalence classes.*'

A notable fact, which we consider one of the main modelling results of our contribution, is that

Fact 2.6.3 *States with the same profile can have different equilibria.* ⊣

Or, more precisely: even when all voters have the same voting preferences, if their knowledge about each other's voting preferences is different, their manipulative behaviour may also be different. We now proceed with the examples.

2.6.1 Examples of conditional equilibria in plurality voting

We reconsider the example of Section 2.3.1 about voters Leela (1) and Sunil (2), voting for an animated movie that may be a (Alice in Wonderland), b (Brave), and c (Cars), where Leela's preferences are $a \succ_1 b \succ_1 c$ and Sunil's preferences are $c \succ_2 b \succ_2 a$. This is the profile P.

As is often the case in game-theoretic models of voting, there exist many Nash equilibria, most of which are irrelevant, i.e, give worse outcome to the voters than their sincere vote. The point of this example however is to demonstrate that equilibria, and thus strategic votes, depend on voters' uncertainty about other voters' preferences, even when their sincere preferences remain the same.

In our running example we now present equilibria when there is: no uncertainty, uncertainty between two profiles, and different kinds of uncertainty between three states (for two profiles). In the following examples, we express the payoffs for both voters by their ranking (0, 1, or 2) for the winner. In the following examples, we use the plurality voting rule, in which only the top-ranked candidate in each ballot is important and we can ignore the rest of the ranking. In other words, here a vote is a single candidate instead of a linear order.

1	2
a	a
b	b
c	c

(a) Profile P

1	2
c	c
b	b
a	a

(b) Profile P'

1/2	a	b	c
a	a	b	a
b	b	b	b
c	a	b	c

(c) Matrix of winners

1/2	a	b	c
a	2.0	[1.1]	2.0
b	1.1	[1.1]	1.1
c	2.0	1.1	0.2

(d) Outcome matrix for P

1/2	a	b	c
a	0.0	[1.1]	0.0
b	[1.1]	[1.1]	1.1
c	0.0	1.1	[2.2]

(e) Outcome matrix for P'

Figure 2.1. Voters have no uncertainty

Example 2.6.4 (No uncertainty) Let P be the profile. If voter 1 votes for a and voter 2 votes for c, then the tie-breaking preferences determines a as the winner, 2's least preferred candidate. A strategic vote of 2 for candidate b, makes b win, a better outcome for voter 2. Equilibrium pairs of votes are (a, b) and (b, b). For voter 1, casting the vote to $a \succ_1 b \succ_1 c$ and $a \succ_1' c \succ_1' b$ is dominant.

The other profile used in the examples in this section is where 1 shares the preferences of 2. This is the profile P'. We also determine the equilibria of that profile. There is no dominant vote. Although (c, c) is an equilibrium vote for P', there are other equilibria that give worse outcomes to the voters compared to their sincere votes.

An overview of the equilibria for P and for P' is in Figure 2.1.Equilibria are boxed. We write $x.y$ instead of (x, y) to denote the payoffs for voters 1 and 2 of the outcome of the election.

Example 2.6.5 (Uncertainty between two profiles) Now consider the profile model consisting of two states t with profile P and u with profile P'. The accessibility relation for voter 1 is the identity on the model and for voter 2 it is the universal relation. Figure 2.2 gives the two-state profile model, the strategic game matrix with conditional votes and winners, and the strategic game matrix with payoffs. Let us explain conditional votes for each voter. First consider voter 1. As voter 1 has two equivalence classes, there are nine conditional votes for her, each of which is denoted by xy, where x and y are her votes in states t and u, respectively. For example, conditional vote ba for 1 means that in state t, 1 votes b and in state u, 1 votes a. Also, vote ba can be interpreted as "If 1 prefers a over c then 1 votes b, and if 1 prefers c over a then 1 votes a."

The voter 2 has one equivalence class, so there are three conditional votes: a, b and c. Conditional profiles are pairs (xy, z) where x is 1's vote in t and y is 1's vote in u, and z is 2's vote in $\{t, u\}$. For example, (ba, a) is a conditional profile that means voter 1 votes for b in state t, and for a in t and voter 2 votes for a in $\{t, u\}$.

The winner matrix (the middle matrix in Figure 2.2) contains pairs xy in which x and y are the winner in state t with profile P respectively in state u with profile P'. For example, for conditional vote (ba, c) we get ba as the entry in the winners matrix: if voter 1 votes b and voter 2 votes c then the tie $(b \succ a \succ c)$ makes b win, and if 1 votes a and 2 votes c then a wins.

The payoff matrix next to the winners matrix contains triples $xy.z$ where x is the rank of the outcome for 1 in state t (P) of the conditional profile, and y is the rank of the outcome for 1 in

(a) Profile models

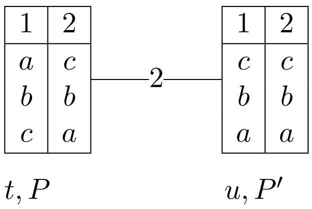

(b) Matrix of winners

1\2	a	b	c
aa	aa	bb	aa
ab	ab	bb	ab
ac	aa	bb	ac
ba	ba	bb	ba
bb	bb	bb	bb
bc	ba	bb	bc
ca	aa	bb	ca
cb	ab	bb	cb
cc	aa	bb	cc

(c) Outcome matrix and equilibria

1\2	a	b	c
aa	20.0	11.1	20.0
ab	21.0	11.1	21.0
ac	20.0	11.1	21.0
ba	10.0	11.1	10.0
bb	11.1	11.1	11.1
bc	10.0	11.1	12.1
ca	20.0	11.1	00.0
cb	21.0	11.1	01.1
cc	20.0	11.1	02.2

Figure 2.2. Voter 2 is uncertain whether voter 1 prefers a over c or c over a

state u (P'), and z is the rank of the worst outcome for 2 in her only equivalence class $\{t, u\}$. For example, for conditional profile (ba, c) we get (10.0) as the entry in the payoff matrix; if 1 votes b and 2 votes c then b wins, so both voters get payoff 1 in P; if 1 votes a and 2 votes c then a wins, they both get payoff 0 in P'; the worst of 0 and 1 is 0, so the payoff for voter 2 of this conditional profile is 0.

Now, let us explain how to determine the equilibria. The equilibria are, maybe, as expected. If the profile is P then it is still dominant for voter 1 to vote a. If the profile is P', voting for c is not dominant for voter 1. Because voter 2 is risk averse, (cc, c) is no longer an equilibrium vote in P'. As 2 is uncertain whether 1 prefers c over a or a over c, the safer (risk avoiding) strategy for 2 is now to vote b, even though 1 and 2 both prefer c. Voter 1 knows this as well.

Voter 2 does not have a dominant vote, because if he assumes that voter 1 always votes c, the best response is also to vote c and not to vote b. So this is the only case where voting b is not an equilibrium vote for 2. ⊣

Example 2.6.6 (Uncertainty between three states) We now add further uncertainty to the two-state profile model where 2 is uncertain between profiles P and P' (states t and u). Consider the following two ways to do this.

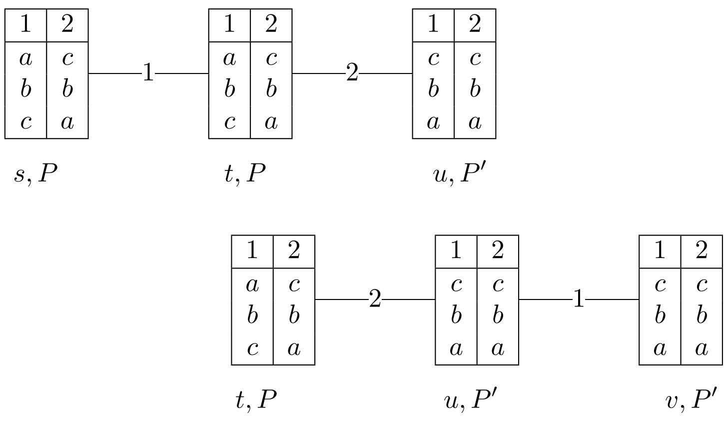

First, consider adding a state s (to t, u) which has the same profile P as the state t. In this three state model voter 2 is not uncertain about P but 1 cannot distinguish s from t. (So the partition for 1 on the domain is $\{s, t\}$, $\{u\}$; and the partition for 2 on the domain is $\{s\}$, $\{t, u\}$.) This is the profile model of Section 2.3.1, where in state s Sunil (2) knows that Leela's (1's) preferences are different from his own, but in state t he does not know that. Does 2 behave differently in s and t? Next, consider adding a state v (to t, u) which has the same profile P' as in state u. In this model, voter 1 but not 2 is uncertain between u and v. In v both 1 and 2 know that they have the same preferences, but in u they don't. Will 2 behave differently in u and v? In both models, voter 1 in all states knows voter 2's preferences. The point of the example is, that it is not rational for 2 to behave (vote) differently in s and in t, but that it is rational for 2 to behave differently in u and in v. Figure 2.3 (on page 27) gives an overview of the conditional equilibria for both profile models, including the matrices with winners in order to calculate the payoffs.

As it may be confusing to see three winners but four payoff values let us explain once more the mechanics of conditional profiles and conditional equilibria. For example, take the t, u, v model. Consider conditional profile (ac, bc), in which we find bbc and 11.12 for the winners matrix entry and the payoff matrix, respectively. Conditional profile (ac, bc) denotes that

- If 1 prefers a (i.e., in state t) then she votes a, and if 1 prefers c (i.e., in states u, v) then she votes c.

- If 2 is uncertain whether 1 prefers a (i.e., in states t, u) then he votes b, and if 2 knows that 1 prefers c (i.e., in state v) then he votes c.

The winners in states t, u, v of these conditional votes are, respectively, b, b, c. If the state is t, then 1 votes a and 2 votes b, so b wins. If u, then 1 votes c and 2 votes b, so b wins. If v, then 1 votes c and 2 votes c, so c wins.

The payoff entry is 11.12 because: for voter 1 in state t, the worst (and only) outcome is b with payoff 1, for voter 1 in states u, v the worst of b and c, with payoffs 1 and 2, is (also) 1; for voter 2 in states t, u the worst of b and b, both with payoff 1, is 1, and for voter 2 the worst and only outcome in state v is c with payoff 2.

Now, we show that conditional profile (ac, bc) is an equilibrium. For this, we have to check four virtual players. For example, player 1 in state t cannot do better, because the first digit of the payoff entries for profiles (bc, bc), (cc, bc) are not greater than 1; player 1 in class $\{u, v\}$ cannot do better: we check the second digit of the entries for conditional profiles (aa, bc) and (ab, bc). Player 2 cannot do better in t, u, check the third digit in entries for (ac, ac) and (ac, cc); and player 2 also cannot do better in v, we check the fourth digit in the entries for (ac, ba) and (ac, bb).

In the s, t, u model, there is no difference in payoffs of voter 2 whether he knows voter 1's preferences. If the profile is P, voter 1 *knows* that voting for a is dominant. On that assumption, voter 2 should vote b, such that b wins. Indeed, in almost all equilibria (except (bc, bc)), b wins and the payoff is 1 for both voters. Unlike for the two-state example, where in all equilibria voter 2 votes b, there are now equilibria in which voter 2 does not vote b. However, these are not really interesting, as 1 votes b in these, which is dominated by 1 voting a.

In the t, u, v model, there is a difference in payoffs of voter 2 depending on whether 1 is uncertain. There are equilibria in which both players vote c in state v, namely (ac, bc) and (bc, bc). Whereas there is no equilibrium in which both players vote c in state u, even though they both prefer c over b and a. We can easily backup this result by our intuitions. If voter 2 is uncertain about 1's preferences (in $\{t, u\}$), the worst-case avoiding strategy remains voting b. If voter 2 knows that 1's preferences are the same as his (in $\{v\}$), even 1's uncertainty is not enough to make him change his vote. The same cannot be said for voter 1. In state v, she has to weigh the odds against voter 2 playing safe and voting b instead of c; but either way, voting c then also

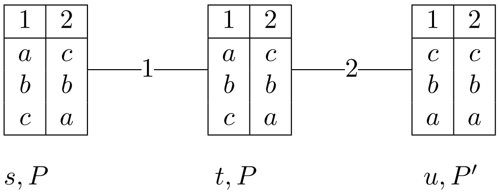

1\2	*aa*	*ab*	*ac*	*ba*	*bb*	*bc*	*ca*	*cb*	*cc*
aa	*aaa*	*abb*	*aaa*	*baa*	*bbb*	*baa*	*aaa*	*abb*	*aaa*
ab	*aab*	*abb*	*aab*	*bab*	*bbb*	*bab*	*aab*	*abb*	*aab*
ac	*aaa*	*abb*	*aac*	*baa*	*bbb*	*bac*	*aaa*	*abb*	*aac*
ba	*bba*	*bbb*	*bba*	*bba*	*bbb*	*bba*	*bba*	*bbb*	*bba*
bb	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*
bc	*bba*	*bbb*	*bbc*	*bba*	*bbb*	*bbc*	*bba*	*bbb*	*bbc*
ca	*aaa*	*abb*	*aca*	*baa*	*bbb*	*bca*	*caa*	*cbb*	*cca*
cb	*aab*	*abb*	*acb*	*bab*	*bbb*	*bcb*	*cab*	*cbb*	*ccb*
cc	*aaa*	*abb*	*acc*	*baa*	*bbb*	*bcc*	*caa*	*cbb*	*ccc*

1\2	*aa*	*ab*	*ac*	*ba*	*bb*	*bc*	*ca*	*cb*	*cc*
aa	20.00	11.01	20.00	10.10	**11.11**	10.10	20.00	11.01	20.00
ab	21.00	11.01	21.00	11.10	**11.11**	11.10	21.00	11.01	21.00
ac	20.00	11.01	22.00	10.10	**11.11**	12.10	20.00	11.01	22.00
ba	10.10	**11.11**	10.10	10.10	**11.11**	10.10	10.10	**11.11**	10.10
bb	11.11	**11.11**	11.11	**11.11**	**11.11**	11.11	11.11	**11.11**	11.11
bc	10.10	**11.11**	12.11	10.10	**11.11**	**12.11**	10.10	**11.11**	12.11
ca	20.00	11.01	00.00	10.10	11.11	10.10	00.20	01.21	00.20
cb	21.00	11.01	01.01	11.10	11.11	01.11	01.20	01.21	01.21
cc	20.00	11.01	02.02	10.10	11.11	02.12	00.20	01.21	02.22

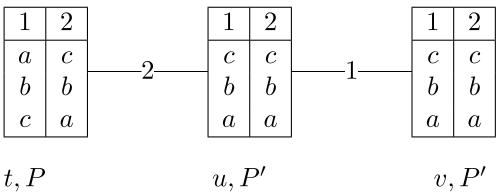

1\2	*aa*	*ab*	*ac*	*ba*	*bb*	*bc*	*ca*	*cb*	*cc*
aa	*aaa*	*aab*	*aaa*	*bba*	*bbb*	*bba*	*aaa*	*aab*	*aaa*
ab	*abb*	*abb*	*abb*	*bbb*	*bbb*	*bbb*	*abb*	*abb*	*abb*
ac	*aaa*	*aab*	*aac*	*bba*	*bbb*	*bbc*	*aca*	*acb*	*acc*
ba	*baa*	*bab*	*baa*	*bba*	*bbb*	*bba*	*baa*	*bab*	*baa*
bb	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*	*bbb*
bc	*baa*	*bab*	*bac*	*bba*	*bbb*	*bbc*	*bca*	*bcb*	*bcc*
ca	*aaa*	*aab*	*aaa*	*bba*	*bbb*	*bba*	*caa*	*cab*	*caa*
cb	*abb*	*abb*	*abb*	*bbb*	*bbb*	*bbb*	*cbb*	*cbb*	*cbb*
cc	*aaa*	*aab*	*aac*	*bba*	*bbb*	*bbc*	*cca*	*ccb*	*ccc*

1\2	*aa*	*ab*	*ac*	*ba*	*bb*	*bc*	*ca*	*cb*	*cc*
aa	20.00	20.01	20.00	10.10	**11.11**	10.10	20.00	20.01	20.00
ab	21.01	21.01	21.01	**11.11**	**11.11**	**11.11**	21.01	21.01	21.01
ac	20.00	20.01	20.02	10.10	11.11	**11.12**	20.00	21.01	22.02
ba	10.00	10.01	10.00	10.10	**11.11**	10.10	10.00	10.01	10.00
bb	11.11	11.11	11.11	**11.11**	**11.11**	**11.11**	11.11	11.11	11.11
bc	10.00	10.01	10.02	10.10	11.11	**11.12**	10.10	11.11	12.12
ca	20.00	20.01	20.00	10.10	**11.11**	10.10	00.00	00.01	00.00
cb	21.01	21.01	21.01	**11.11**	**11.11**	**11.11**	01.11	01.11	01.11
cc	20.00	20.01	20.02	10.10	11.11	11.12	00.20	01.21	02.22

Figure 2.3. Conditional equilibria for profile models where two states have the same profile

gives her best result. So for voter 1 voting b and c give the same outcome considering the worst case scenario, and this is indeed the case: (ab, bc) and (bb, bc) are also equilibria. ⊣

2.7 Revealing voting preferences

We can extend the setting for the interaction of voting preferences and knowledge of the previous sections with operations in which voters are informed of other's preferences, thus reducing ('updating') their uncertainty. An obvious choice for such updates is the *public announcement* [135] of propositions about profiles, such as voters informing each other about their preferences.

The dynamics of public announcements can be modelled as an operation $(\mathcal{M}, s) \mapsto ((\mathcal{M}|T), s)$, where $T \subseteq S$ is the denotation in $\mathcal{M}$ of a proposition about profiles φ, and $\mathcal{M}|T$ means model restriction to subdomain T. In that case we also write $(\mathcal{M}|\varphi)_s$.

Given a knowledge profile $(\mathcal{M}, s)$, the precondition for execution of the operation *public announcement of* φ (or *update with* φ) is that φ is true in $(\mathcal{M}, s)$, and the way to execute it is to restrict the model $\mathcal{M}$ to all the states where φ is true. We can then investigate the truth of propositions about profiles in that model restriction. This therefore allows us to evaluate more complex propositions about profiles, namely of shape 'after update with φ, ψ (is true),' such as: 'After voter 1 reveals her preferences to voter 2, voter 2 knows that she has a successful manipulation' (but before that update, she didn't know). This is embodied in the following definitions, see Section 2.8 for a formal logical setting.

Definition 2.7.1 (Updated knowledge profile) *Let $(\mathcal{M}, s)$ be a knowledge profile, where $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$, and let φ be a proposition about profiles with denotation $S' \subseteq S$ and such that $s \in S'$. Then the* updated knowledge profile *$((\mathcal{M}|\varphi), s)$ is defined as $\mathcal{M}|\varphi = (S', \{\sim'_1, \ldots, \sim'_n\}, V')$ where $\sim'_i = \sim_i \cap (S' \times S')$ and for all $P \in O(\mathcal{C})^n$, $V' = V|_{S'}$.*

Let ψ also be a proposition about profiles. In $(\mathcal{M}, s)$ the proposition ψ is true after update with φ, if whenever φ is true in $(\mathcal{M}, s)$, ψ is true in $((\mathcal{M}|\varphi), s)$. ⊣

It requires some reflection to realize what can count as a public announcement in the setting of voting. Truthful public announcement logic, and other dynamic epistemic logics, are logics of observation and do not have a notion of agency. The idea is that information coming from a further unnamed external source is considered reliable enough to incorporate it (or merely to investigate the consequences of incorporating it, for example when planning future actions). The semantics of 'public announcement' merely reflects the effect on the agent's uncertainties of the result of incorporating information. Anything else, such as reasons to incorporate it, or doubts about the source or the provenance of the information, are not modelled. The further important assumption is that the information is public, i.e., that all agents observe the information similarly. This has a certain legal connotation, as in proclaiming laws: agents cannot be excused from *not* having incorporated the information, by, say, not having paid attention to what was announced. It is not that this cannot be modelled: one can also model defective channels, and also announcements of lies; but those are extensions of the framework from which public announcement logic is an abstraction. Unfortunately the terminology does not help to put the reader on the right foot: both the 'truthful' and the 'announcement' in truthful public announcement logic are misnomers. There are other types of observation than listening (to what is said), such as seeing; and given the anonymous source, no distinction can be made between 'true' and 'truthful' — the former would therefore have been more appropriate. While there is no agency, it is common in dynamic epistemic logic to import agency by the backdoor, namely by modeling an announcement φ by an agent (voter) a as the public announcement of $K_a\varphi$ (agent a knows φ). This backdoor is also a trapdoor: if φ were not announced by the agent, but $K_a\varphi$ by an external source, that would result in the same information change.

This background on the meaning of 'announcement' may help to justify that what voters reveal about their own preferences can be modelled as truthful public announcements. A voter revealing her preferences only to another voter does not make a public but a *private* announcement, that is not observed by all voters (unless of course there are only two voters, as in the examples that we give below). This requires a more complex framework. A voter revealing her preferences to me is indistinguishable from someone else revealing to me that voter's preferences. Lack of sincerity when revealing preferences is not considered.

We proceed with some obvious results on the preservation of knowledge and manipulation after updates, followed by less obvious results on the preservation and emergence of conditional equilibria after updates, and concluded (in the next subsection) by a detailed example.

Proposition 2.7.2 *Successful manipulation is preserved after update. Equilibrium profiles are preserved after update.* ⊣

Proof The existence of a successful manipulation depends only on the profile of the actual state, the point s of the knowledge profile $(\mathcal{M}, s)$. This remains the actual state after any update. Similarly, for equilibrium profiles. □

Proposition 2.7.3 *Knowledge of successful manipulation is preserved after update.* ⊣

Proof Let $(\mathcal{M}, s)$ be a knowledge profile such that $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$. Let φ be a proposition with denotation $S' \subseteq S$, and let $(\mathcal{M}|\varphi, s)$ be the updated knowledge profile such that $\mathcal{M}|\varphi = (S', \{\sim'_1, \ldots, \sim'_n\}, V')$. Suppose that voter i knows 'de dicto' that she can successfully manipulate $(\mathcal{M}, s)$. Then, by Def. 2.4.1, it means for all $t \in S$ such that $s \sim_i t$ she can successfully manipulate $V(t)$. Since $[s]_{\sim_i} \cap S' \subseteq [s]_{\sim_i} \subseteq S$, Prop. 2.7.2 implies that voter i knows 'de dicto' that she can successfully manipulate $(\mathcal{M}|\varphi, s)$. The proof for 'de re' knowledge is similar. □

Proposition 2.7.4 *Weak knowledge of successful manipulation is not preserved after update.* ⊣

Proof For weak knowledge of manipulation there were two requirements: (a) the profile of at least one state in the equivalence class for voter i has a successful manipulation, and (b) that manipulation gives an equal or better outcome for the profiles of all states in that equivalence class. The state with the manipulation need not be the actual state, therefore, after model restriction the existential requirement (a) may no longer hold. (A counterexample can be easily constructed.) □

It is on first sight less clear what happens to conditional equilibria after updates. We first have to define the update of a conditional profile.

Definition 2.7.5 (Updated conditional profile) *Let profile model* $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$ *and conditional profile* $[P] = \{[\succ]_1, \ldots, [\succ]_n\}$ *be given, where* $[\succ]_i : S\backslash\sim_i \to O(\mathcal{C})$ *are conditional votes. Now let* φ *be a proposition about profiles such that* $\mathcal{M}|\varphi = (S', \{\sim'_1, \ldots, \sim'_n\}, V')$*. Then the conditional profile* $[P']$ *after the update with* φ *consists of conditional votes* $[\succ']_i : S'\backslash\sim'_i \to O(\mathcal{C})$ *defined as: for all* $s \in S'$, $[\succ']_i([s]_{\sim'_i}) = [\succ]_i([s]_{\sim_i})$. ⊣

What may happen to a conditional vote $[\succ]_i : S\backslash\sim_i \to O(\mathcal{C})$ is that (1) an equivalence class for i disappears, namely if none of the states in that class satisfies the update φ, (2) an equivalence class for i shrinks, because some states satisfy φ and others do not, and (3) an equivalence class for i remains the same, because all of its states satisfy φ. In the first case, the virtual voter (in the Bayesian game sense) $(i, [s]_{\sim_i})$ ceases to exist; the conditional vote has one less condition. In the second case we have that $[s]_{\sim'_i} \subset [s]_{\sim_i}$. This may affect the payoff for i of vote $\succ_i$ in that class, because states with the least value may have been removed (namely if $\min\{F(V(s)) \mid s \in [s]_{\sim_i}\} \prec_i \min\{F(V'(s)) \mid s \in [s]_{\sim'_i}\}$). In the third case there is no difference in payoff.

Proposition 2.7.6 *Given a conditional profile that is an equilibrium, the updated conditional profile may not be an equilibrium.* ⊣

Proof The proposition says, in other words, that *conditional equilibrium is not preserved after update.* We recall that a conditional profile is an equilibrium if no agent has a successful manipulation in any of its equivalence classes, i.e., if the Bayesian game played between all virtual voters $(i, [s]_{\sim_i})$, where $[s]_{\sim_i}$ is an equivalence class for voter i, has the corresponding profile as an equilibrium.

As said above, three different things can happen to a conditional profile after an update: equivalence classes may disappear, so that virtual voters disappear from the game; equivalence classes remain the same, so that the worst-case payoff of that class also remains the same; or equivalence classes may become smaller, so that the worst outcome of the vote in that class may no longer be so bad as before.

This may affect an equilibrium as follows. The payoff for a virtual voter $(i, [s]_{\sim_i})$ casting vote $\succ'_i$ in the conditional equilibrium is the worst outcome for voter i in class $[s]_{\sim_i}$. This payoff is greater or equal to (at least as good as) the worst outcome in class $[s]_{\sim_i}$ for any other vote $\succ''_i$. This is fragile and not preserved after update: Namely, if a $t \in [s]_{\sim_i} \backslash [s]_{\sim'_i}$ exists such that

$$F(V(t)) \quad \succ_i \quad F(V(t)_{-i}, \succ''_i)$$

then we have that

$$\begin{array}{rlcl} & \min\{F(V(t)) \mid t \in [s]_{\sim_i}\} & \succeq_i & \min\{F(V(t)_{-i}, \succ''_i) \mid t \in [s]_{\sim_i}\} \\ \text{whereas} & \min\{F(V'(t)) \mid t \in [s]_{\sim'_i}\} & \prec_i & \min\{F(V'(t)_{-i}, \succ''_i) \end{array}$$

where as before V' and $\sim'_i$ are the profile function and accessibility relation in the updated profile model; Section 2.7.1 provides an example. □

Proposition 2.7.7 *Given a conditional profile that is not an equilibrium, the updated conditional profile may be an equilibrium.* ⊣

Proof In other words, **not** *being a conditional equilibrium is also not preserved after update.* This requires that $\min\{F(V(t)) \mid t \in [s]_{\sim_i}\} \prec_i \min\{F(V(t)_{-i}, \succ''_i) \mid t \in [s]_{\sim_i}\}$ and $\min\{F(V'(t)) \mid t \in [s]_{\sim'_i}\} \succeq_i \min\{F(V'(t)_{-i}, \succ''_i)\}$. Again, the next subsection provides an example, for the plurality voting rule. □

So conditional equilibria can both disappear and appear after updates. It is yet unclear to us if there are general patterns here. However, additional strategic behaviour comes into the picture with these negative results. An update may consist of a voter revealing her voting preferences. It may be that this voter's sincere vote is not part of an equilibrium conditional profile, but that after this voter reveals her sincere vote, the updated conditional profile is an equilibrium. Or, for another example, a voter may know that after revealing (part of) her voting preferences, other voters may now vote differently because other voters' preferences are now part of an equilibrium. This interaction between strategic voting and strategic communication may well open up new vistas in social choice theory. We will summarily address this as well in the next example section.

2.7.1 Examples of updates in plurality voting

Consider again the examples for plurality voting of Section 2.3.1 and Section 2.6.1. The two-state profile model $\mathcal{M}$ of Section 2.6.1 is reprinted again below, in the middle, together with two updates, one on the left and one on the right.

In $\mathcal{M}$ in state t, after voter 1 informs voter 2 of her true preference $a \succ_1 c$, no uncertainty remains, and 1 and 2 commonly know that the actual profile is P. We recall that equilibrium votes for P are $(a, b), (b, b)$. All conditional equilibria are preserved after update. Conditional equilibria for the t, u profile model had shape (xy, b), where x is voter 1's vote in t and y is voter 1's vote in u, and $xy \neq cc$. For example, given conditional profile (bc, b), the updated conditional profile according to Definition 2.7.5 is (b, b). There is therefore no strategic incentive for voter 1 to inform voter 2. Because if she informs voter 2 about her preference, then 2 votes b, and the outcome is not better than before for 1.

On the other hand, in state u voter 1 has an incentive to make her preferences known to 2. In the model with P and P', there is no equilibrium in which 2 votes c. But after 1 informs 2 that her preferences are the same as his preferences, (c, c) is an equilibrium. (And, as both voters then vote sincerely, this equilibrium is of more strategic value than the other equilibria.) Most conditional equilibria of the two-state model are preserved, but not those where 1 votes c. For example, (ac, b) was an equilibrium, but the updated profile (c, b) is not an equilibrium.

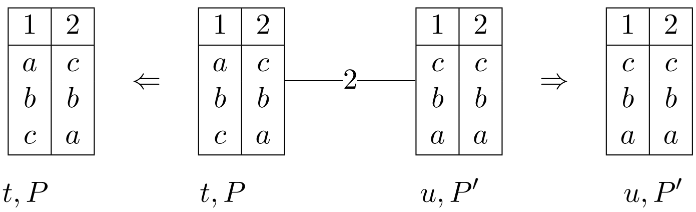

In principle we could give a similar story for the three-state profile models, but we have already made our point: in this example voter 1 has a strategic interest to reveal her preferences to voter 2 because subsequently, she can expect the outcome of the vote to be better than before. Before update it was b, after update it was c, and $c \succ_1 b$. In other words, *when there is uncertainty about votes, voters have different ways of acting strategically: voting strategically or strategically revealing voting preferences.*

2.8 A logic of knowledge and voting

In the present section, we present a logic of knowledge and voting.

2.8.1 Syntax and semantics

Given n voters $Ag = \{1, \dots, n\}$, m candidates $\mathcal{C} = \{a, b, \dots\}$, profiles $O(\mathcal{C})^n$, the language $\mathcal{L}_{KV}$ is basically the language of public announcement logic [49] with a special set of atomic propositions $\mathsf{At} = O(\mathcal{C})^n \cup \mathcal{C}$.

Definition 2.8.1 (Logical language) *The formulas of the language $\mathcal{L}_{KV}$ are defined by the following grammar:*

$$\varphi ::= P \mid x \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid [\varphi]\varphi$$

where $i \in Ag$, $x \in \mathcal{C}$, and $P \in O(\mathcal{C})^n$. We use all of the standard abbreviations for disjunction ($\vee$), implication ($\rightarrow$) and bi-implication ($\leftrightarrow$). ⊣

A propositional variable x stands for 'the winner is x' (the interaction with a given profile is prescribed with an axiom **F** given later); $K_i\varphi$ stands for 'voter i knows that φ'; $[\varphi]\psi$ stands for 'after public announcement of φ, ψ (is true)'.

We define other voting concepts by notational abbreviation. The vote of voter i is defined as the disjunction of all profiles containing it. Similarly we define that voter i prefers one candidate over another (computational efficiency is not our goal):

$$\begin{array}{llll} \text{For all linear orders } \succ_i \in O(\mathcal{C})^n: & \succ_i & := & \bigvee_{\{P \mid P_i = \succ_i\}} P \\ \text{For all linear orders } \succ_i \in O(\mathcal{C})^n: & a \succ_i b & := & \bigvee_{\{P \mid a \succ_i b \text{ for } P_i = \succ_i\}} \end{array}$$

We therefore liberally allow overloading in order to keep the further logical formalisation readable. In view of readability we may enclose expressions of form $\succ_i'$ or $a \succ_i' b$ between parentheses.

Now, we define models, and the satisfaction relation for the language $\mathcal{L}_{KV}$.

Definition 2.8.2 ($\mathcal{L}_{KV}$-model) *Given a set of n voters $Ag = \{1, \ldots, n\}$, a set of m candidates $\mathcal{C}$, an $\mathcal{L}_{KV}$-*model *is a structure $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$, where S is a set of states, for all all $i \in Ag$, $\sim_i$ is an equivalence relation on S, and $V\colon \mathsf{At} \to \mathcal{P}(S)$ is a valuation in which $\mathsf{At} = O(\mathcal{C})^n \cup \mathcal{C}$.* ⊣

The next definition introduces a class of $\mathcal{L}_{KV}$-models that corresponds to profile models of Def. 2.3.1.

Definition 2.8.3 (Profile $\mathcal{L}_{KV}$-model) *An $\mathcal{L}_{KV}$-model $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\})$ is a* profile $\mathcal{L}_{KV}$-model, *if*

1. *Each state has exactly one profile assigned to it, that is, for all $P, P' \in O(\mathcal{C})^n$: if $P \neq P'$ then $V(P) \cap V(P') = \emptyset$, and for all $s \in S$ there is $P \in O(\mathcal{C})^n$ such that $s \in V(P)$, and*

2. *all voters know their own preferences, that is, for all $i \in Ag$, for all $P \in O(\mathcal{C})^n$ and for all $s, s' \in S$: $s \sim_i s'$ implies ($s \in V(P)$ iff $s' \in V(P)$).* ⊣

So far we treated the voting function F as a meta-level concept, it was not part of the definition of the $\mathcal{L}_{KV}$-models. Now, we deal with it more explicitly to introduce a class of profile $\mathcal{L}_{KV}$-models in which a given voting rule F determines the winner candidate.

Definition 2.8.4 (F-voting model) *Let $F\colon O(\mathcal{C})^n \to \mathcal{C}$ be a voting rule. An $\mathcal{L}_{KV}$-model $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$ is an F-*voting model *if it is a profile $\mathcal{L}_{KV}$-model and*

1. *for $x, x' \in \mathcal{C}$: if $x \neq x'$ then $V(x) \cap V(x') = \emptyset$, and*

2. *for all $P \in O(\mathcal{C})^n$, for all $s \in S$: $s \in V(P)$ implies $s \in V(F(P))$.*

Item (1) says that there is at most one winner in each state, and item (2) says that if P is the profile at state s, then the winner at s is $F(P)$. ⊣

The $\mathcal{L}_{KV}$-formulas are interpreted on pointed $\mathcal{L}_{KV}$-models $(\mathcal{M}, s)$ consisting of an $\mathcal{L}_{KV}$-model $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$ and a state $s \in S$.

Definition 2.8.5 (Semantics) *Given an $\mathcal{L}_{KV}$-model $\mathcal{M} = (S, \{\sim_1, \ldots, \sim_n\}, V)$, we define when formula $\varphi \in \mathcal{L}_{KV}$ is true in $(\mathcal{M}, s)$, also written as $(\mathcal{M}, s) \models \varphi$, as follows:*

$$\begin{array}{lll}
(\mathcal{M}, s) \models P & \textit{iff} & s \in V(P) \\
(\mathcal{M}, s) \models x & \textit{iff} & s \in V(x) \\
(\mathcal{M}, s) \models \neg\varphi & \textit{iff} & (\mathcal{M}, s) \not\models \varphi \\
(\mathcal{M}, s) \models \varphi \wedge \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ and } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models K_i\varphi & \textit{iff} & \textit{for every } t \textit{ such that } s \sim_i t, (\mathcal{M}, t) \models \varphi \\
(\mathcal{M}, s) \models [\varphi]\psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ implies } ((\mathcal{M}|\varphi), s) \models \psi
\end{array}$$

where $\mathcal{M}|\varphi = (S', \{\sim'_1, \ldots, \sim'_n\}, V')$ with $S' = \{t \in S : (\mathcal{M}, t) \models \varphi\}$, $\sim'_i = \sim_i \cap (S' \times S')$, $V'(s) = V(s) \cap S'$, for all $s \in S'$. The expression $(\mathcal{M}, s) \not\models \varphi$ stands for 'not $((\mathcal{M}, s) \models \varphi)$'. If $(\mathcal{M}, s) \models \varphi$ for all $s \in S$, we write $\mathcal{M} \models \varphi$ (φ is valid on $\mathcal{M}$*) and if this is the case for all $\mathcal{M}$, we say that φ is* valid*, and we write $\models \varphi$.* ⊣

We now present principles that are valid on the class of profile models, and that will feature as axioms in the proof system. Let $\overline{\bigvee}$ denote the exclusive disjunction which is a binary operator that expresses a disjunction in which exactly one of the formulas is true.

$$\begin{array}{lll}
\mathbf{P} & : & \overline{\bigvee}_{P \in O(\mathcal{C})^n} P \\
\mathbf{C} & : & \overline{\bigvee}_{x \in \mathcal{C}} x \\
\mathbf{VF}_F & : & \bigwedge_{P \in O(\mathcal{C})^n} (P \to F(P)) \qquad F \text{ is a voting rule.} \\
\mathbf{N} & : & \bigwedge_{i \in Ag} \bigwedge_{\succ_i \in O(\mathcal{C})} ((\succ_i) \to K_i(\succ_i))
\end{array}$$

Axiom **P** and **C** spell out, respectively, that there is exactly one profile is assigned to a state, and that there is exactly one winner of the election. Axiom **F** is the definition of the voting function F in the logic. The final axiom **N** says that voters know their own (complete) preferences.

Proposition 2.8.6 *Let F be a voting rule, and let $\mathcal{M}$ an $\mathcal{L}_{KV}$-model. The axioms* **P**, **C**, $\mathbf{VF}_F$, *and* **N** *are valid on $\mathcal{M}$ iff $\mathcal{M}$ is a F-voting model.*

Proof Straightforward. □

The logic of knowledge and voting is now defined essentially as public announcement logic [135] over F-voting models.

Definition 2.8.7 (The logic of knowledge and voting) *Given a voting rule F,* the logic of knowledge and F-voting *is the set of $\mathcal{L}_{KV}$-formulas that are valid in all F-voting models.* ⊣

Proposition 2.8.8 *Let F be a voting rule. The logic of knowledge and F-voting has a complete axiomatization over the class of F-voting models.* ⊣

Proof The axiomatization of public announcement logic is standard [135]. To this we add the axioms **P**, **C**, $\mathbf{VF}_F$, and **N**. For the completeness we observe that the canonical model of the logic without public announcements is an F-voting model (Prop. 2.8.6), and that the completeness of the logic with announcements is as usual obtained because every formula is equivalent to one without announcements (the axioms are rewriting rules, pushing all logical connectives beyond announcements, such as $[\varphi](\psi \wedge \chi) \leftrightarrow ([\varphi]\psi \wedge [\varphi]\chi)$). □

2.8.2 Example

We succinctly demonstrate the logic using the running example of Section 2.3.1 (page 15) about Leela and Sunil, where the knowledge profile is $(\mathcal{M}, t)$.

- Leela prefers Alice in Wonderland over Cars:
  $(\mathcal{M}, t) \models a \succ_1 c$

- Sunil does not know that Leela prefers Alice in Wonderland over Cars:
  $(\mathcal{M}, t) \models \neg K_2(a \succ_1 c)$

- Leela knows what Sunil likes, but Leela is uncertain whether Sunil knows what Leela likes (i.e., she is uncertain whether Sunil knows Leela's preferences):
  $(\mathcal{M}, t) \models K_1(\succ_2) \wedge \neg K_1(K_2(\succ_1) \vee K_2\neg(\succ_1)) \wedge \neg K_1\neg(K_2(\succ_1) \vee K_2\neg(\succ_1))$

- Leela and Sunil know their own preferences (on the entire model):
  $\mathcal{M} \models ((\succ_1) \rightarrow K_1(\succ_1)) \wedge ((\succ_2) \rightarrow K_2(\succ_2))$

- Sunil does not know that Leela prefers Alice in Wonderland over Cars, but after he was told so, he knows it:
  $(\mathcal{M}, t) \models \neg K_2(a \succ_1 c) \wedge [a \succ_1 c]K_2(a \succ_1 c)$

2.9 Conclusion

Our results We presented a logic for the interaction of voting and knowledge. The semantic primitive is the knowledge profile: a profile including uncertainty of voters about what the actual profile is. This reveals different notions of knowledge of manipulation. In particular, we introduced 'de re' knowledge of manipulation and 'de dicto' knowledge of manipulation, and novel notions of equilibria, including conditional equilibrium for risk-averse voters. We further modelled the dynamics of uncertainty such as revealing preferences (by voters or by a central authority) and its effects on knowledge of manipulation and conditional equilibrium. We proved that knowledge of successful manipulation is preserved after such updates but that conditional equilibria are not.

Declaring votes by variable assignments Another form of dynamics than that of revealing voters' preferences is the dynamics of *declaring votes*. Just as there may be uncertainty about truthful votes, there may also be uncertainty about declared votes. This is relevant for the investigation of *safe manipulation* [150], where the manipulating voter announces her vote to a (presumably large) set of voters sharing her preferences but is unsure of how many will follow her, and also in Stackelberg voting games, in which voters declare their votes in sequence, following a fixed, exogenously defined order.

Revealing preferences is informational ('purely epistemic') change, whereas declaring votes is ontic/factual change. A dynamic epistemic logic equivalent of a declared vote is a so-called *public assignment* [48, 20]. A succinct way to expand our framework with uncertainty about declared votes is to add a duplicate set of propositional variables for voter preferences, to represent their declared votes. Initially setting all these variables to false, declaring a vote then becomes an assignment setting such a variable to true (whereas all other possible votes, for that voter, remain false).

Central authority Apart from the n voters, it is convenient to distinguish yet another agent: a designated agent who is the *central authority*, or *chair*. This opens the door to the logical modelling of well-studied problems in computational social choice, such as control by the chair, or determining possible winners. In this work, the chair manifests as tie-breaking rule, and the reason we did not explicitly model the chair is that her role is uniform throughout a knowledge profile model. We assume that there is no uncertainty on what the voting rule (and thus the tie-breaking preferences) is, we only considered uncertainty about preferences of other voters. So in that sense the chair is exogenous.

The chair as designated agent can be called agent 0 (thus distinguishing the chair from the voters $1, \ldots, n$). The universal relation on a knowledge profile model can then be seen as the indistinguishability relation of the agent 0, the central authority. On a connected model (i.e., when there is always a path (with respect to the indistinguishability relation) between any two states in the model) this is the same as common knowledge of the voters about their uncertainty. The computational tasks of the central authority, such as determining the possible winners or other form of control, may well be harder on knowledge profiles as it has to take uncertainty into account. Identifying the central authority with an agent with universal access allows us to determine whether it is harder.

Coalitional manipulation Coalitions play a big role in voting, because in realistic settings the power of individual voters is very limited. Coalitional notions also play an important role in epistemic logic. Two notions useful in our setting are common knowledge and distributed knowledge. Given a knowledge profile, a proposition is commonly known in a coalition G, if it is true in all states reachable (from the actual state of the knowledge profile) by arbitrarily long finite paths in the model (reflexive transitive closure of the union of all accessibility relations for voters in G). With the interpretation of common knowledge of coalition G we can associate an equivalence relation $\sim_G$, defined as $(\bigcup_{i \in G} \sim_i)^*$. A proposition is distributedly known in a knowledge profile, if it is true in states reachable from the actual state via the intersection of accessibility relations, i.e., the relation $\bigcap_{i \in G} \sim_i$.

By analogy, just as the vote of an individual agent depends on her knowledge, the vote of a coalition would seem to depend on the common knowledge of that coalition. But that seems wrong. In voting theory, the power of a coalition means the power of a set of agents that can decide on a joint action *as a result of communication between them.* Communication makes the uncertainty about each others' profiles disappear. In terms of knowledge profiles, this means that we are talking about another model, namely the model where for all agents $i \in G$, $\sim_i$ is refined to $\bigcap_{i \in G} \sim_i$. What determines the voting power of a coalition is not common knowledge of that coalition but distributed knowledge of that coalition, and involves an update of the knowledge profile model. That is possible, but makes for an unlucky marriage of modelling constraints. A more suitable restriction seems only to consider coalitions of voters having the same uncertainty (i.e., $\bigcap_{i \in G} \sim_i = \sim_i$ for all $i \in G$). Having the same uncertainty determines a *type* of voter. That makes sense in voting and is a common modelling constraint: we only consider coalitions of the same *type*. One can then define knowledge of manipulation and conditional equilibria for coalitions.

Knowledge and belief We modelled *knowledge* of preferences. We did not model *belief* of preferences. Unlike knowledge, beliefs may be incorrect. Somewhat similarly, unlike truthful announcements, insincere announcements (e.g., lying about your true preferences in a voting poll) may lead to false beliefs.

Consider the following variant in which two voters are uncertain between profiles P and P': voter 1 knows which of P and P' is the case, but voter 2 believes (incorrectly) that P is the case, i.e., te actual state is u.

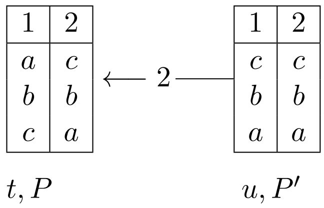

In state u of this knowledge profile, voter 2 will now not vote c, because he believes that voter 1 prefers a, to which b is the best response (with plurality voting, and tie breaking rule $b \succ a \succ c$). Therefore, he will vote b. Unlike Example 2.6.5, he will do that even if he is not risk-averse.

Changing from knowledge to belief allows for truly counterintuitive scenarios, such as agents believing their preferences to be different from what they really are. For example, above, swap the preferences of 1 and 2 in P; i.e., $c \succ_1 b \succ_1 a$ and $a \succ_2 b \succ_2 c$. We now have that if voter 2 really prefers c, then he believes that he prefers a.

Clearly, the interaction between belief and voting is more complex than between knowledge and voting. Technically there are few issues, the same logical framework as in Section 2.8 can be used with minor adjustments.

Applications The logical setting defined in the chapter allows us to represent various classes of situations already studied specifically in (computational) social choice, thus offering a general representation framework in which, of course, new classes of problems will be representable as well, thus providing an homogeneous, unified representation framework. To represent such classes of problems we need the extensions of the framework that were discussed above: uncertainty for coalitions, explicit modelling of the chair, and assignments (to represent declaring votes) instead of merely announcements (to represent revealing preference). As a final example, we mention the issue of *possible and necessary winners* [100]: Let there be one more agent (the chair), who has incomplete knowledge of the votes. Voter x is a possible winner if the chair does not know that x is not a (co)winner, and a necessary winner if the chair knows that x is a (co)winner. Describing such knowledge and ignorance and its strategic consequences would be a typical application of our framework.

3

Algebraic semantics of refinement action logic

Contents

3.1 Introduction

In *modal logic* we attempt to formalise propositions about *possibility* and *necessity*. In epistemic modal logics the modal operator is interpreted as knowledge or belief [94], initially for a single knowing agent but later for a set of agents, including their higher-order knowledge (i.e., what they know about each other) [58]. The knowledge of agents is encoded in a relational structure known as a *Kripke model* or *relational structure*, consisting of a domain of worlds, a binary accessibility relation for each agent, and a valuation of atomic propositions over the worlds. Informative updates can be formalised as yet another modal operator, a dynamic modality, that is interpreted as a relation between such Kripke models. A well-known form of informative updates are action models [11], wherein the updates themselves also take the shape of a relational structure.

The Kripke model resulting from executing an action model in an initial Kripke model can also be seen as a *refinement* of that initial model. A refinement relation is like a bisimulation relation, except that from the three relational requirements only 'Atoms' and 'Zag' need to be satisfied. This therefore results in structural loss. From the perspective of knowledge change, this implies that in the refined model agents know more, namely they are less uncertain between different worlds. In [31] refinement modal logic (**RML**) is introduced, wherein modal logic is augmented with a new operator $\exists$ (and with its dual $\forall$, they are interdefinable as usual), which quantifies over all refinements of a given pointed model. In this logic the expression $\exists\varphi$ stands for "there is a refinement after which φ." In other words, $\exists\varphi$ is true in a Kripke model $\mathcal{M}$ with point s (we write $(\mathcal{M}, s)$ for such a pair) if there is a pointed model $(\mathcal{M}', s')$ such that $((\mathcal{M}, s), (\mathcal{M}, s)')$ is a pair in the refinement relation, (we also say that $(\mathcal{M}', s')$ is a refinement of $(\mathcal{M}, s)$), and such that φ is true in $(\mathcal{M}', s')$. The logic is equally expressive as basic modal logic [31]. A well-known result is that action model execution results in a refinement [45]. We can similarly (although not trivially) augment the modal logic of knowledge with refinement quantifiers, and also the multi-agent logic of knowledge [31].

A different form of quantification is over action models. This has been investigated in [84]. This logic is called arbitrary action model logic [84]. It is an extension of action model logic with an action model quantifier where $\exists\varphi$ stands for "there is an action model such that after its execution φ (is true)." Given such an expression $\exists\varphi$, in [84] Hales presents a method for synthesizing a multi-pointed action model α_T after which φ is true (in the sense that $\exists\varphi$ is logically equivalent to $\langle\alpha_T\rangle\varphi$), and he also proved that the action model quantifier is equivalent to the refinement quantifier. To show this equivalence, he introduced *refinement action model logic* (**RAML**), that extends action model logic (**AML**) with refinement quantifiers. The syntax and semantics of **RAML** are formed by combining the syntax and semantics of **AML** and **RML**.

In this chapter we develop an algebraic semantics of refinement action modal logic (**RAML**). Already from close to the inception of dynamic epistemic logics, there has been a strong current to model such logics in algebraic or coalgebraic settings [9, 10]. More recently, in [105, 114] an algebraic semantics was proposed for public announcement logic and action model logic. This methodology has further been productively used in [42] for a probabilistic dynamic epistemic logic and in Chapter 4 of this thesis for epistemic updates on bilattices.

In [105, 114], product updates are dually characterised through a construction that transforms the complex algebra associated with a given Kripke model into the complex algebra associated with the model updated by means of an action model. Given a Kripke model $\mathcal{M}$ and an action model α, the result of executing that action model can be seen as a submodel of a so-called intermediate model that contains copies of $\mathcal{M}$ indexed by the domain of α. In this way, action model logic can be endowed with an algebraic semantics that is dual (and equivalent) to the relational one, via a Jónsson-Tarski-type duality [25]. In particular, this holds for the multi-pointed action model α_T such that $\exists\varphi$ is equivalent to $\langle\alpha_T\rangle\varphi$, according [84] mentioned above.

We use this result to define the algebraic semantics of **RAML**. Indeed, we can dually characterise the algebraic notion of refinement relation as a lax-morphism (named *refinement morphism*) between the complex algebras associated with a given initial Kripke model and a 'resulting' Kripke model that is in the refinement relation with the initial model. Then, via the Jónsson-Tarski duality, we associate that resulting Kripke model to a modal Boolean algebra. Given the set of all refinements of the initial Kripke model, we then take the product of all corresponding modal Boolean algebra in order to define a unique algebra and the required refinement morphism. The motivation behind our approach is to capture the non-constructive notion of refinement. Whereas arbitrary action model logic approaches the notion of refinement with brute force by having a witnessing action model that enforces the same postcondition φ

bound by the quantifier, refinement modal logic only needs the existence of such an epistemic action (and thus the possibility of synthesizing it) but not the actual construction.

Overview. In Section 3.2, we recall basic definition and results on modal logic, refinement modal logic, action model logic, and refinement action model logic. In Section 3.3, we recall relevant algebraic concepts and fix notation. In Section 3.4.1, we present the methodology based on [105, 114] to define the algebraic semantics of dynamic epistemic logics. Finally, in Section 3.5, we present the algebraic semantics of refinement action modal logic. Section 3.6 describes our results in view of prior works and concludes.

3.2 Logical preliminaries

In this section, we briefly recall basic modal logic [25], action model logic [11], refinement modal logic [45, 31], and refinement action model logic [85]. Throughout this chapter, we assume a countable set of atomic propositions At.

To ease the presentation, we present here the single-agent version of these logics. All results in this section generalize to the multi-agent setting.

3.2.1 Basic modal logic

We recall standard definitions from modal logic. For an extensive introduction to modal logic we refer to [25].

Definition 3.2.1 (Language $\mathcal{L}_\Box$) *The language of basic modal logic is inductively defined as:*

$$\mathcal{L}_\Box \ni \varphi ::= p \in \mathsf{At} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi$$

We use the standard abbreviations from propositional logic. We also employ the abbreviation $\Diamond\varphi ::= \neg\Box\neg\varphi$. ⊣

The formulas of $\mathcal{L}_\Box$ are interpreted over relational structures known as Kripke models [101, 94].

Definition 3.2.2 (Kripke frame/model) *A Kripke frame $\mathcal{F} = \langle S, R\rangle$ is a pair where S is the* domain *consisting of* worlds *(or* states*), and $R \subseteq S \times S$ is a binary accessibility relation. A* Kripke model*, or in this chapter just* model*, is a triple $\mathcal{M} = \langle S, R, V\rangle$ where $\langle S, R\rangle$ is a Kripke frame and $V : \mathsf{At} \to \mathcal{P}(S)$ is a* valuation *assigning to each atomic proposition $p \in \mathsf{At}$ the subset of the domain where the atomic proposition p is true.* ⊣

For every Kripke frame $\mathcal{F} = \langle S, R\rangle$ and every $s \in S$, we write $R(s)$ to denote the set of successors of s, and we write $t \in R(s)$ to denote that $(s, t) \in R$.

Given $s \in S$, a *pointed Kripke frame* is a pair $(\mathcal{F}, s)$, written as $\mathcal{F}_s$. A *multi-pointed Kripke frame* is a pair $(\mathcal{F}, T)$, written as $\mathcal{F}_T$, consists of a Kripke frame $\mathcal{F}$ along with a non-empty set of designated states $T \subseteq S$. is a *multi-pointed frame* denoted $\mathcal{F}_T$. Similarly, given $s \in S$, and a non-empty set $T \subseteq S$, a pair $(\mathcal{M}, s)$ is a *pointed model*, and $(\mathcal{M}, T)$ is a multi-pointed Kripke model. We now define the semantics of modal logic.

Definition 3.2.3 (Semantics of basic modal logic) *Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model. The interpretation of $\varphi \in \mathcal{L}_\Box$ is defined inductively by*

$$\begin{array}{lll}
(\mathcal{M}, s) \models p & \textit{iff} & s \in V(p) \\
(\mathcal{M}, s) \models \varphi \wedge \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ and } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \neg\varphi & \textit{iff} & (\mathcal{M}, s) \not\models \varphi \\
(\mathcal{M}, s) \models \Box\varphi & \textit{iff} & \textit{for all } t \in R(s) : (\mathcal{M}, t) \models \varphi
\end{array}$$

We write $[\![\varphi]\!]_{\mathcal{M}}$ *to denote the set of states satisfies* φ, *where* $[\![\varphi]\!]_{\mathcal{M}} = \{s \in S : (\mathcal{M}, s) \models \varphi\}$. ⊣

Let $\mathcal{M} = \langle S, R, V\rangle$ be a pointed Kripke model and let $\varphi \in \mathcal{L}_\Box$. The formula φ is *valid* on $\mathcal{M}$, notation $\mathcal{M} \models_{\mathbf{K}} \varphi$, if for all $s \in S$, $(\mathcal{M}, s) \models \varphi$. If $\mathcal{M} \models_{\mathbf{K}} \varphi$ for every Kripke model $\mathcal{M}$ then we say that φ is valid and write $\models_{\mathbf{K}} \varphi$. The basic modal logic (**K**) is defined as the set of validities of $\mathcal{L}_\Box$ over the class of all Kripke models.

In the following sections, we will use the *cover modality* ∇ following the definitions of [24]. For any finite set $\Phi \subseteq \mathcal{L}_\Box$ of formulas we define $\nabla\Phi$ as the syntactic abbreviation

$$\nabla\Phi := \Box \bigvee_{\varphi\in\Phi} \varphi \wedge \bigwedge_{\varphi\in\Phi} \Diamond\varphi.$$

where $\bigvee_{\varphi\in\emptyset} \varphi := \mathsf{f}$ (always false) and $\bigwedge_{\varphi\in\emptyset} \varphi := \mathsf{t}$ (always true). We next recall *cover disjunctive normal form.*

Definition 3.2.4 (Cover disjunctive normal form) *The set of formulas in* $\mathcal{L}_\Box$ *that are in* cover disjunctive normal *form is generated by the following grammar:*

$$\varphi ::= \pi \wedge \nabla\Phi \mid \varphi \vee \varphi$$

where π *is a propositional formula, and* Φ *is a finite set of formulas in cover disjunctive normal form.* ⊣

Lemma 3.2.5 *[cf. [83, Lemma 5.1.2]] Every formula* $\varphi \in \mathcal{L}_\Box$ *is equivalent to a formula in cover disjunctive normal form in the logic* **K**. ⊣

Definition 3.2.6 (Bisimulation) *Let* $\mathcal{M} = \langle S, R, V\rangle$ *and* $\mathcal{M}' = \langle S', R', V'\rangle$ *be Kripke models. A binary relation* $Z \subseteq S \times S'$ *is a (Kripke)* bisimulation *between* $\mathcal{M}$ *and* $\mathcal{M}'$ *if for all* $(s, s') \in Z$*:*

(Atoms) s *and* s' *satisfy the same atomic propositions;*

(Zig) *for all* $t \in S$, *if* $t \in R(s)$, *then there exists* $t' \in S'$ *such that* $t' \in R'(s')$ *and* $(t, t') \in Z$*;*

(Zag) *for all* $t' \in S'$, *if* $t' \in R'(s')$, *then there exists* $t \in S$ *such that* $t \in R(s)$ *and* $(t, t') \in Z$.

We write $(\mathcal{M}, s) \leftrightarroweq (\mathcal{M}', s')$ *(and say that* $(\mathcal{M}, s)$ *and* $(\mathcal{M}', s')$ *are bisimilar) iff there exists a bisimulation between* $\mathcal{M}$ *and* $\mathcal{M}'$ *that links* s *and* s'. ⊣

We now recall the axiomatisation for basic modal logic **K**.

Definition 3.2.7 (Axiomatisation $\mathbb{K}$) *The axiomatisation* $\mathbb{K}$ *is a substitution schema consisting of the following axioms and rules:*

P	*All propositional tautologies*
K	$\Box(\varphi \leftrightarrow \psi) \leftrightarrow (\Box\varphi \leftrightarrow \Box\psi)$
MP	*From* $\varphi \leftrightarrow \psi$ *and* φ *infer* ψ
NecK	*From* φ *infer* $\Box\varphi$

where $\varphi, \psi \in \mathcal{L}_\Box$. ⊣

3.2.2 Action model logic

This section recalls definitions and results from the action model logic of Baltag, Moss and Solecki [11]. For an extensive introduction on action model logic we refer the reader to [49].

Definition 3.2.8 (Language $\mathcal{L}_{\Box\alpha}$) *Let* At *be as countable set of atomic propositions. The set $\mathcal{L}_{\Box\alpha}$ of formulas and the set $\mathcal{AM}_{AML}$ of action models are defined by mutual induction by the following grammar:*

$$\mathcal{L}_{\Box\alpha} \ni \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid [\alpha_k]\varphi$$
$$\mathcal{AM}_{AML} \ni \alpha ::= ((K, R_\alpha), (\varphi_1, ...\varphi_n))$$

where $p \in \mathsf{At}$, $\alpha \in \mathcal{AM}_{AML}$, k is a state in α, (K, R_α) is a finite Kripke frame, and $\varphi_1, ..., \varphi_n \in \mathcal{L}_{\Box\alpha}$, where $n = |K|$. ⊣

The state space K of an action model $((K, R), \varphi_1, ..., \varphi_n)$ is understood to be enumerated as $K = \{1, ..., n\}$ such that the sequence $\varphi_1, ..., \varphi_n$ corresponds to a function $Pre_\alpha : K \to \mathcal{L}_{\Box\alpha}$ which we call the *precondition function*. From now on we write $\langle K, R_\alpha, Pre_\alpha\rangle$ for action models, where $Pre_\alpha : K \to \mathcal{L}_{\Box\alpha}$. The notation α_k is formally to be understood as the pair (α, k), which we call a *pointed action model* or *epistemic action*. A *multi-pointed action model* is a pair $\alpha_T = \langle\alpha, T\rangle$ where $\alpha \in \mathcal{AM}_{AML}$ is an action model and $T \subseteq K$.

We assume all the standard abbreviations from modal logic, in addition to the abbreviations

$$\langle\alpha_k\rangle\varphi := \neg[\alpha_k]\neg\varphi, \quad \langle\alpha_T\rangle\varphi := \bigvee_{k\in T}\langle\alpha_k\rangle\varphi, \quad [\alpha_T]\varphi := \bigwedge_{k\in T}[\alpha_k]\varphi,$$

where $T \subseteq K$.

Definition 3.2.9 (Semantics of action model logic) *Let $\mathcal{M} = (S, R, V)$ be a Kripke model and let s be a state in $\mathcal{M}$. The interpretation of $\varphi \in \mathcal{L}_{\Box\alpha}$ on the pointed Kripke model $(\mathcal{M}, s)$ is the same as its interpretation in modal logic, defined in Definition 3.2.3, with an additional inductive case:*

$$(\mathcal{M}, s) \models [\alpha_k]\varphi \quad \textit{iff} \quad (\mathcal{M}, s) \models Pre_\alpha(k) \textit{ implies } (\mathcal{M}_\alpha, (s, k)) \models \varphi$$

where $\mathcal{M}_\alpha = \langle S_\times, R_\times, V_\times\rangle$ is the product update *defined as*

$$S_\times = \{(t, j) \in S \times K \mid (\mathcal{M}, t) \models Pre_\alpha(j)\}$$
$$(s, k)R_\times(s', k') \quad \textit{iff} \quad (t, t') \in R \;\textit{ and }\; (j, j') \in R_\alpha$$
$$(t, j) \in V_\times(p) \quad \textit{iff} \quad t \in V(p)$$

We will also say that $(\mathcal{M}_\alpha, (s, k))$ is the result of executing α_k in the pointed model $(\mathcal{M}, s)$. We write $[\![\varphi]\!]_{\mathcal{M}}$ to denote the set of states in a Kripke model $\mathcal{M}$ that satisfy φ. A formula $\varphi \in \mathcal{L}_{\Box\alpha}$ is valid *on $\mathcal{M}$, notation $\mathcal{M} \models_{\mathbf{AML}} \varphi$, if for all $s \in S$, $(\mathcal{M}, s) \models \varphi$. If $\mathcal{M} \models_{\mathbf{AML}} \varphi$ for all Kripke models $\mathcal{M}$ then we say φ is valid and write $\models_{\mathbf{AML}} \varphi$.* ⊣

The logic **AML** is the set of $\mathcal{L}_{\Box\alpha}$ formulas that are valid.

Definition 3.2.10 (Axiomatisation $\mathbb{AML}$) *The axiomatisation $\mathbb{AML}$ [84, Definition IV.1] consists of the rules and axioms of $\mathbb{K}$ along with the following axioms and the rule of necessitation for dynamic box modalities:*

AP	$[\alpha_k]p \leftrightarrow (Pre_\alpha(k) \to p)$
AN	$[\alpha_k]\neg\varphi \leftrightarrow (Pre_\alpha(k) \to \neg[\alpha_k]\varphi)$
AC	$[\alpha_k](\varphi \wedge \psi) \leftrightarrow ([\alpha_k]\varphi \wedge [\alpha_k]\psi)$
AK	$[\alpha_k]\Box\varphi \leftrightarrow (Pre_\alpha(k) \to \bigwedge\{\Box[\alpha_{k'}]\varphi \mid k' \in R_\alpha(k)\})$
NecA	*From* φ *infer* $[\alpha_k]\varphi$

where $\varphi, \psi \in \mathcal{L}_{\Box\alpha}$, $\alpha \in \mathcal{AM}_{AML}$, *and* $p \in \mathsf{At}$. ⊣

Given a formula $\varphi \in \mathcal{L}_{\Box\alpha}$, $\vdash_{\mathbb{AML}} \varphi$ means that φ is a theorem in $\mathbb{AML}$, meaning that one can derive φ from the axioms and rules of $\mathbb{AML}$. The axiomatisation $\mathbb{AML}$ is sound and complete with respect to the semantics of the logic **AML** [49].

3.2.3 Refinement modal logic

In this subsection, we recall the basic definitions and results on refinement modal logic of [31]. We begin with the definition of the refinement relation between Kripke models. Recall from Def. 3.2.6 that a bisimulation between two Kripke models is a binary relation that satisfies three conditions: **(Atoms)**, **(Zig)**, and **(Zag)**. A refinement is like a bisimulation except it only needs to satisfy conditions **(Atoms)** and **(Zag)**.

Definition 3.2.11 (Refinement) *Let $\mathcal{M} = \langle S, R, V\rangle$ and $\mathcal{M}' = \langle S', R', V'\rangle$ be two Kripke models. A binary relation $\mathfrak{R} \subseteq S \times S'$ is a* refinement *between $\mathcal{M}$ and $\mathcal{M}'$ if for all $(s, s') \in \mathfrak{R}$:*

(Atoms) *s and s' satisfy the same atomic propositions;*

(Zag) *for all $t' \in S'$, if $t' \in R'(s')$, then there exists $t \in S$ such that $t \in R(s)$ and $(t, t') \in \mathfrak{R}$.*

We write $(\mathcal{M}, s) \succeq (\mathcal{M}', s')$ is there exists a refinement $\mathfrak{R}$ between $\mathcal{M}$ and $\mathcal{M}'$ with $(s, s') \in \mathfrak{R}$. ⊣

It is clear that every bisimulation between two Kripke models is a refinement between them [31].

Now, we recall the syntax and semantics of the refinement modal logics.

Definition 3.2.12 (Language of refinement modal logic) *The* language of refinement modal logic $\mathcal{L}_{\Box\forall}$, *is inductively defined as:*

$$\mathcal{L}_{\Box\forall} \ni \varphi ::= p \in \mathsf{At} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid \forall\varphi$$

We use all of the standard abbreviations from modal logic, in addition to the abbreviations $\exists\varphi := \neg\forall\neg\varphi$. ⊣

Definition 3.2.13 (Semantics of refinement modal logic) *Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model and let s be a state in $\mathcal{M}$. The interpretation of a formula $\varphi \in \mathcal{L}_{\Box\forall}$ is the same as its interpretation in modal logic, defined in Definition 3.2.3, with an additional inductive case:*

$$(\mathcal{M}, s) \models \forall\varphi \quad \textit{iff} \quad \textit{for all } (\mathcal{M}', s') : \textit{if } (\mathcal{M}, s) \succeq (\mathcal{M}', s') \textit{ then } (\mathcal{M}', s') \models \varphi.$$

Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model and let $\varphi \in \mathcal{L}_{\Box\forall}$. The formula φ is valid on $\mathcal{M}$, notation $\mathcal{M} \models_{\mathbf{RML}} \varphi$, if for all $s \in S$, $(\mathcal{M}, s) \models \varphi$. If $\mathcal{M} \models_{\mathbf{RML}} \varphi$ for every Kripke model $\mathcal{M}$ then we say φ is valid and write $\models_{\mathbf{RML}} \varphi$.

Definition 3.2.14 (Axiomatisation $\mathbb{RML}$) *The axiomatisation $\mathbb{RML}$ [31] consists of the rules and axioms of $\mathbb{K}$ along with the following axioms and the rules:*

$$\begin{array}{rl} \textbf{R} & \forall(\varphi \to \psi) \to \forall\varphi \to \forall\psi \\ \textbf{RProp} & \forall p \leftrightarrow p \textit{ and } \forall\neg p \leftrightarrow \neg p \\ \textbf{RK} & \exists\nabla\Phi \leftrightarrow \bigwedge \Diamond\exists\Phi \\ \textbf{NecR} & \textit{From } \varphi \textit{ infer } \forall\varphi \end{array}$$

where $\varphi, \psi \in \mathcal{L}_{\Box\forall}$, Φ is a finite subset of $\mathcal{L}_{\Box\forall}$, and $p \in \mathsf{At}$. ⊣

Given a formula $\varphi \in \mathcal{L}_{\Box\forall}$, $\vdash_{\mathbb{RML}} \varphi$ means that φ is a theorem in $\mathbb{RML}$, meaning that one can derive φ from the axioms and rules of $\mathbb{RML}$.

We recall an important result connecting refinement modal logic to action modal logic that we will use later.

Lemma 3.2.15 *[45, Prop. 4& 5] The result of executing an epistemic action in a pointed model is a refinement of that model. Dually, for every refinement of a* finite *pointed model there is an epistemic action such that its execution results in a model bisimilar to that refinement.* ⊣

We note that there is no result that shows the refinement of an infinite pointed model is the result of the execution of an epistemic action.

3.2.4 Refinement action model logic

In this subsection, we recall the syntax and semantics of refinement action model logic of [85]. The results we recall here will play a vital role in section 3.5.

Definition 3.2.16 (Language $\mathcal{L}_{\Box\alpha\forall}$) *Let At be a countable set of atomic propositions. The set $\mathcal{L}_{\Box\alpha\forall}$ of formulas and the set $\mathcal{AM}_{RAML}$ of action models are defined by mutual induction by the following grammar:*

$$\begin{aligned} \mathcal{L}_{\Box\alpha\forall} \ni \varphi &::= p \in \mathsf{At} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi \mid [\alpha_k]\varphi \mid \forall\varphi \\ \mathcal{AM}_{RAML} \ni \alpha &::= ((K, R_\alpha), (\varphi_1, ...\varphi_n)) \end{aligned}$$

where $\alpha \in \mathcal{AM}_{RAML}$, k is a state in α, (K, R_α) is a finite Kripke frame, and $\varphi_1, ..., \varphi_n \in \mathcal{L}_{\Box\alpha\forall}$, where $n = |K|$. We use all of the standard abbreviations from action model logic and refinement modal logic. ⊣

Definition 3.2.17 (Semantics of $\mathcal{L}_{\Box\alpha\forall}$) *Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model and let s be a state in $\mathcal{M}$. The interpretation of a formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ is the same as its interpretation in action modal logic and refinement modal logic.* ⊣

Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model and let $\varphi \in \mathcal{L}_{\Box\alpha\forall}$. The formula φ is valid on $\mathcal{M}$, notation $\mathcal{M} \models_{\mathbf{RAML}} \varphi$, if for all $s \in S$, $(\mathcal{M}, s) \models \varphi$. If $\mathcal{M} \models_{\mathbf{RAML}} \varphi$ for every Kripke model $\mathcal{M}$ then we say φ is valid and write $\models_{\mathbf{RAML}} \varphi$. The logic **RAML** is defined as the set of validities of $\mathcal{L}_{\Box\alpha\forall}$.

The logic **RAML** agrees with **AML** and **RML** on formulas from their respective sublanguages because the syntax and semantics of **RAML** are formed by combining the semantics of **AML** and **RML**. Moreover, a sound and complete axiomatisation for **RAML** can be given by combining the axiomatisations $\mathbb{AML}$ and $\mathbb{RML}$

Definition 3.2.18 (Axiomatisation $\mathbb{RAML}$) *The axiomatisation $\mathbb{RAML}$ [84] is a substitution schema consisting of the rules and axioms of $\mathbb{AML}$ and $\mathbb{RML}$.* ⊣

For a given formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, we write $\vdash_{\mathbb{RAML}} \varphi$, if φ is derivable in $\mathbb{RAML}$.

Hales in [84] shows that if there exists a refinement in which a given formula φ is satisfied then one can construct a finite action model that results in φ being satisfied. He provides an algorithm that computes such a finite multi-pointed action model for any formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$.

Theorem 3.2.19 *[84, Theorem V.3] Let $\varphi \in \mathcal{L}_{\Box\alpha\forall}$. Then there exists a multi-pointed action model α^{φ}_{T} such that $\vdash_{\mathbb{RAML}} [\alpha^{\varphi}_{T}]\varphi$ and $\vdash_{\mathbb{RAML}} \langle\alpha^{\varphi}_{T}\rangle\varphi \leftrightarrow \exists\varphi$.* ⊣

In Section 3.5, we will use the actual construction of this result to prove soundness with respect to algebraic semantics. We therefore describe here the inductive construction of the action model for atomic propositions and formulas of the form $\nabla\Phi$.

1. For a given atomic proposition p, we construct the multi-pointed action model $\alpha^{p}_{K^p} = ((K^p, R^p, Pre^p), K^p)$ as follows:

$$\begin{aligned} K^p &= \{k^\star, \mathsf{skip}\} \\ R^p_\alpha &= \{(k^\star, \mathsf{skip}), (\mathsf{skip}, \mathsf{skip})\} \\ Pre^p_\alpha &= \{(k^\star, p), (\mathsf{skip}, \mathsf{t})\} \end{aligned}$$

   Given a pointed Kripke model $(\mathcal{M}, s)$, we have

$$(\mathcal{M}, s) \models \langle\alpha^{p}_{K^p}\rangle p \quad \text{iff} \quad (\mathcal{M}, s) \models p.$$

   Then since $(\mathcal{M}, s) \models p$ iff $(\mathcal{M}, s) \models \exists p$ [31, Theorem 28], it follows that

$$(\mathcal{M}, s) \models \langle\alpha^{p}_{K^p}\rangle p \quad \text{iff} \quad (\mathcal{M}, s) \models \exists p. \tag{3.1}$$

2. Let Φ be a finite set of $\mathcal{L}_{\Box\alpha\forall}$-formulas. For every $\varphi \in \Phi$ let $\alpha^{\varphi}_{K^\varphi} = ((K^\varphi, R^\varphi_\alpha, Pre^\varphi_\alpha), K^\varphi)$ be a multi-pointed action model such that $\vdash_{\mathbb{RAML}} [\alpha^{\varphi}_{K^\varphi}]\varphi$ and $\vdash_{\mathbb{RAML}} \langle\alpha^{\varphi}_{K^\varphi}\rangle\varphi \leftrightarrow \exists\varphi$. Without loss of generality we can assume that each of the K^φ are pair-wise disjoint. Then we construct the action model $\alpha^{\nabla\Phi}_{K^{\nabla\Phi}} = ((K^{\nabla\Phi}, R^{\nabla\Phi}_\alpha, Pre^{\nabla\Phi}_\alpha), K^{\nabla\Phi})$ as follows:

$$\begin{aligned} K^{\nabla\Phi} &= \{k^\star\} \cup \bigcup_{\varphi\in\Phi} K^\varphi \\ R^{\nabla\Phi}_\alpha &= \{(k^\star, k) \;:\; \varphi \in \Phi, k \in K^\varphi\} \cup \bigcup_{\varphi\in\Phi} R^\varphi_\alpha \\ Pre^{\nabla\Phi}_\alpha &= \{(k^\star, \bigwedge_{\varphi\in\Phi} \Diamond\exists\varphi)\} \cup \bigcup_{\varphi\in\Phi} Pre^\varphi_\alpha. \end{aligned}$$

   As [84] has pointed out, $\alpha^{\nabla\Phi}$ is formed by the disjoint union of each action model α^φ, and no outward edges are added to any state K^φ in $\alpha^{\nabla\Phi}$. Then we have

$$\vdash_{\mathbb{RAML}} [\alpha^{\nabla\Phi}_{K^{\nabla\varphi}}]\nabla\varphi \quad \text{and} \quad \vdash_{\mathbb{RAML}} \exists\nabla\Phi \leftrightarrow \langle\alpha^{\nabla\varphi}_{K^{\nabla\varphi}}\rangle\nabla\Phi \tag{3.2}$$

   Hence, given a pointed Kripke model $(\mathcal{M}, s)$, we have $(\mathcal{M}, s) \models \exists\nabla\Phi$ iff $(\mathcal{M}, s) \models \langle\alpha^{\nabla\Phi}_{K^{\nabla\Phi}}\rangle\nabla\Phi$.

We note that the action model synthesis described by Hales in [84] heavily relies on the following fact:

Fact 3.2.20 *The logics* **RAML**, **RML**, **AML** *and* **K** *are all equally expressive [11, 31, 84]* .⊣

3.3 Algebraic preliminaries

The aim of the present section is to collect relevant definitions and results on (modal) Boolean algebras. For a more thorough introduction to universal algebra, we refer to [34].

Boolean algebras. A *Boolean algebra* $\mathbf{A} = \langle A, \vee, \wedge, \neg, \mathbf{0}, \mathbf{1} \rangle$ is an algebra with two binary operations $\vee$ (called 'join' or 'or') and $\wedge$ (called 'meet' or 'and'), one unary operation $\neg$ (called 'not' or 'complement'), and two nullary operations $\mathbf{0}$ and $\mathbf{1}$ (called 'bottom' and top') which satisfy the following equations:

$$
\begin{array}{lll}
x \wedge y = y \wedge x & x \vee y = y \vee x & \text{(commutativity)} \\
x \vee (y \vee z) = (x \vee y) \vee z & x \wedge (y \wedge z) = (x \wedge y) \wedge z & \text{(associativity)} \\
x \wedge (x \vee y) = x & x \vee (x \wedge y) = x & \text{(absorption)} \\
x \wedge \mathbf{1} = x & x \vee \mathbf{0} = x & \text{(identity)} \\
x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z) & x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z) & \text{(distributivity)} \\
x \wedge \neg x = \mathbf{0} & x \vee \neg x = \mathbf{1} & \text{(complementation)}
\end{array}
$$

We call the class of Boolean algebras **BA**. We may label the operators and constants with the name of the algebra, as $\vee_{\mathbf{A}}$, $\mathbf{0}_{\mathbf{A}}$, etc., to distinguish them from those in other algebras.

Let X be a set and $\mathcal{P}(X)$ be the set of all the subsets of X. Denote with $\cup$, $\cap$ and $(-)^c$ the operations union, intersection and complement on $\mathcal{P}(X)$, respectively. Then, the structure $\langle \mathcal{P}(X), \cup, \cap, (-)^c, \emptyset, X \rangle$ forms a Boolean algebra, so-called *powerset algebra*.

Underlying poset of a Boolean algebra. A Boolean algebra $\mathbf{A} = \langle A, \vee, \wedge, \neg, \mathbf{0}, \mathbf{1} \rangle$ can also be seen as a *partially ordered set* (poset). A poset is a pair $(A, \leq)$ where A is a set and $\leq$ is a *reflexive*, *antisymmetric* (if $x \leq y$, $y \leq x$ then $x = y$) and *transitive* relation on A. Given a Boolean algebra $\mathbf{A}$, the relation $\leq$ can be defined as follows: for every $x, y \in A$

$$x \leq y \quad \text{iff} \quad x \wedge y = x \quad \text{iff} \quad x \vee y = y$$

We call $(A, \leq)$ the underlying poset of $\mathbf{A}$. Let $(A, \leq)$ be a poset, $x \in A$ and $S \subseteq A$, x is an *upper bound* (resp. *lower bound*) of S, if $s \leq x$ (resp. $x \leq s$) for every $s \in S$. The element $x \in A$ is the *least upper bound* (lub) of S if it is an upper bound of S and if $x \leq s$ for every upper bound u of S. The element $x \in A$ is the *greatest lower bound* (glb) of S if it is a lower bound of S and if $s \leq x$ for every lower bound s of S. If they exist, the least upper bound of S is denoted by $\bigvee S$ and the greatest lower bound of S by $\bigwedge S$. For any Boolean algebra $\mathbf{A} = \langle A, \vee, \wedge, \neg, \mathbf{0}, \mathbf{1} \rangle$, $\bigvee S$ and $\bigwedge S$ of a finite subset $S \subseteq A$ always exist and are unique, however they may not exist if S is infinite. This leads us to the following definition:

Complete Boolean algebra A Boolean algebra $\mathbf{A} = \langle A, \vee, \wedge, \neg, \mathbf{0}, \mathbf{1} \rangle$ is *complete* if $\bigvee S$ and $\bigwedge S$ exist for every $S \subseteq A$.

For an arbitrary set X, the powerset algebra $\langle \mathcal{P}(X), \cup, \cap, (-)^c, \emptyset, X \rangle$ is complete. The underlying order is given by the inclusion $\subseteq$. From a logical perspective, the existence of arbitrary joins and meets enables us to reason about properties of infinite sets.

Definition 3.3.1 (Modal Boolean algebra) *A* modal Boolean algebra *is a structure* $\mathbb{A} = \langle A, \vee, \wedge, \neg, \{\Diamond_i\}_{i \in I}, \mathbf{0}, \mathbf{1} \rangle$ *where* $\langle A, \wedge, \vee, \neg, \mathbf{0}, \mathbf{1} \rangle$ *is a Boolean algebra, I is a non-empty finite set, and* $\Diamond_i : \mathbb{A} \to \mathbb{A}$ *for every* $i \in I$*. A* normal *modal Boolean algebra is a modal Boolean algebra* $\mathbb{A}$ *such that the unary operations* $\{\Diamond_i\}_{i \in I}$ *on* $\mathbb{A}$ *satisfying the following conditions: for every* $i \in I$ *and* $x, y \in A$

1. *(Normality)* $\Diamond_i \mathbf{0} = \mathbf{0}$, *and*
2. *(Additivity)* $\Diamond_i(x \vee y) = \Diamond_i x \vee \Diamond_i y$. $\dashv$

Convention 3.3.2 *Throughout the present chapter, and unless specified otherwise, we assume modal Boolean algebras are normal with only one operator* $\Diamond$*, and we define a dual operator* $\Box$ *of* $\Diamond$ *by* $\Box x := \neg\Diamond\neg x$*. We denote by* **MBA** *the class of normal modal Boolean algebras. For the sake of smooth presentation, we will sometimes write* $\langle \mathbb{A}, \Diamond \rangle$ *instead of* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \mathbf{0}, \mathbf{1}, \Diamond \rangle$.$\dashv$

We now define the standard notion of structure preserving maps between algebras.

Definition 3.3.3 (MBA-Homomorphism) *Let* $\mathbb{A}$ *and* $\mathbb{A}'$ *be two modal Boolean algebras. A map* $f : A \to A'$ *is an* **MBA***-homomorphism, notation* $f : \mathbb{A} \to \mathbb{A}'$*, if for all* $x, y \in A$

$$\begin{aligned} & f(x \wedge_{\mathbb{A}} y) = f(x) \wedge_{\mathbb{A}'} f(y) \\ & f(x \vee_{\mathbb{A}} y) = f(x) \wedge_{\mathbb{A}'} f(y) \\ & f(\neg_{\mathbb{A}} x) = \neg_{\mathbb{A}'} f(x) \\ & f(\Diamond_{\mathbb{A}} x) = \Diamond_{\mathbb{A}'} f(x) \\ & f(c_{\mathbb{A}}) = c_{\mathbb{A}'}. \qquad (c \in \{\mathbf{0}, \mathbf{1}\}) \end{aligned}$$

An **MBA**-isomorphism is a bijective **MBA**-homomorphism and we say two modal Boolean algebra $\mathbb{A}$ and $\mathbb{A}'$ are *isomorphic*, if there is an **MBA**-isomorphism between them (notation: $\mathbb{A} \cong \mathbb{A}'$).

Every **MBA**-homomorphism between two modal Boolean algebra $\mathbb{A}$ and $\mathbb{A}'$ is a *monotone* map between their underlying posets. Recall that a map $f : (A, \leq_A) \to (A', \leq'_A)$ between two posets is *monotone* if whenever $x \leq_A y$ then $f(x) \leq'_A f(y)$ for all $x, y \in A$.

A closely related notion to **MBA**-homomorphism that we will use further on is the notion of *congruence*. A *congruence* θ on a modal Boolean algebra $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond, \mathbf{0}, \mathbf{1} \rangle$ is an equivalence relation on A satisfying the following compatibility property: for all $x, x', y, y' \in A$, if $x \; \theta \; y$ and $x' \; \theta \; y'$ then:

$$(\neg x) \; \theta \; (\neg y), \qquad (x \vee x') \; \theta \; (y \vee y'), \qquad (x \wedge x') \; \theta \; (y \wedge y') \qquad \textit{and} \qquad (\Diamond x) \; \theta \; (\Diamond y).$$

Congruences provides us with a way to construct *quotient algebras*. Let $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond, \mathbf{0}, \mathbf{1} \rangle$ be a modal Boolean algebra, and let θ be a congruence on $\mathbb{A}$. The quotient algebra $\mathbb{A}/\theta$ has as its carrier A/θ, the set of the equivalence classes defined by the congruence θ, namely $A/\theta = \{[x]_\theta \; : \; x \in A\}$ with $[x]_\theta = \{x' \in A \; : \; x \; \theta \; y\}$. There is a natural way to define the operations $\vee', \wedge', \neg'$ and $\Diamond'$ on the set A/θ of equivalence classes of $\mathbb{A}$ over θ. Namely, for all $x, y \in A$, we define

$$\begin{aligned} [x]_\theta \vee' [y]_\theta &:= [x \vee y]_\theta, \\ [x]_\theta \wedge' [y]_\theta &:= [x \wedge y]_\theta, \\ \neg'[x]_\theta &:= [\neg x]_\theta \\ \Diamond'[x]_\theta &:= [\Diamond x]_\theta. \end{aligned}$$

It can be easily shown that the operators defined above, are well-defined, in particular $\Diamond$ is a normal modality. Hence, the quotient algebra $\mathbb{A}/\theta := \langle A/\theta, \vee', \wedge', \neg', \Diamond', [\mathbf{0}]_\theta, [\mathbf{1}]_\theta \rangle$ is a modal Boolean algebra.

As we will see in Section 3.4.1, a congruence θ on the $\Diamond$-free reduct of a modal Boolean algebra $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond, \mathbf{0}, \mathbf{1}\rangle$ is not necessarily a congruence with respect to $\Diamond$. In this situation, one may need to modify the definition of the modality on A/θ to obtain a modal Boolean algebra, so-called *pseudo-quotient* algebra.

The other way to form new algebras is to make a big algebra out of a collection of small ones.

Products. Let $(\mathbb{A}_i)_{i\in I}$ be a family of modal Boolean algebras where $\mathbb{A}_i = \langle A_i, \wedge_i, \vee_i, \neg_i, \Diamond_i, \mathbf{0}_i, \mathbf{1}_i\rangle$ and I is an index set. We define the *product* $\prod_{i\in I}\mathbb{A}_i$ as the modal Boolean algebra $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond_{\mathbb{B}}, \mathbf{0}, \mathbf{1}\rangle$, where A is the Cartesian product $\prod_{i\in I} A_i$ that is defined as

$$\prod_{i\in I} A_i := \{(x_i)_{i\in I} \ : \ \forall i \in I, x_i \in A_i\}$$

with the canonical projections $\pi_j : \prod_{i\in I} A_i \to A_j$ defined as $\pi_j((x_i)_{i\in I}) := x_j$, and the operators are defined cordinatewise; that is, for elements $(x_i)_{i\in I}, (y_i)_{i\in I} \in \prod_{i\in I} A_i$,

$$\begin{aligned} (x_i)_{i\in I} \wedge (y_i)_{i\in I} &:= (x_i \wedge_i y_i)_{i\in I} \\ (x_i)_{i\in I} \vee (y_i)_{i\in I} &:= (x_i \vee_i y_i)_{i\in I} && (\bullet \in \{\vee, \wedge\}) \\ \neg(x_i)_{i\in I} &:= (\neg_i x_i)_{i\in I} \\ \Diamond(x_i)_{i\in I} &:= (\Diamond_i x_i)_{i\in I} \\ c &:= (c_i)_{i\in I}. && (c \in \{\mathbf{0}, \mathbf{1}\}) \end{aligned}$$

We now recall the important concept over posets called *adjunction* that provides us with a way to produce a complete Boolean algebra [44, 34].

Definition 3.3.4 (Adjunction) *A pair (f, g) of monotone maps $f : A \to A'$ and $g : A' \to A$ between two posets $(A, \leq_A)$ and $(A', \leq'_A)$ forms an* adjunction *(notation: $f \dashv g$) if $f(x) \leq'_A y$ is equivalent to $x \leq_A g(y)$, for all $x \in A$ and $y \in A'$. If $f \dashv g$, then g is a* right adjoint *and f a* left adjoint. $\dashv$

The importance of adjunctions stems from the interaction between adjoint maps and joins and meets [44]. Let (f, g) be such that $f \dashv g$, then f preserves existing arbitrary joins if and only if g preserves existing arbitrary meets.

As modalities $\Diamond$ and $\Box$ can be seen as monotone operators on the underlying poset of a modal Boolean algebra, it is natural to ask whether they have adjoint modalities? The answer is: yes. In fact, there exist two unary operators $\blacklozenge$ and $\blacksquare$ such that $\Diamond \dashv \blacksquare$ and $\blacklozenge \dashv \Box$. The modalities $\blacklozenge$ and $\blacksquare$ can be thought of as the backward looking diamond and forward looking box of tense logic, respectively. Every modal Boolean algebra $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond_{\mathbb{B}}, \mathbf{0}, \mathbf{1}\rangle$ expanded with $\blacksquare$ is called a *tense* modal Boolean algebra (notation: $\langle \mathbb{A}, \blacksquare\rangle$). **MBA** provides as an algebraic semantics of modal logic. The following class of powerset algebras plays a key role in algebraizing modal logic.

Complex algebras. Let $\mathcal{F} := \langle S, R\rangle$ be a Kripke frame. The *complex algebra* of $\mathcal{F}$ (notation: $\mathcal{F}^+$) is the powerset algebra $\langle \mathcal{P}(S), \cup, \cap, (-)^c, \emptyset, S\rangle$ expanded with the modality $\Diamond_R : \mathcal{P}(S) \to \mathcal{P}(S)$ defined as

$$\Diamond_R(X) := \{s \in S : t \in R(s) \ \text{ for some } \ t \in X\} = R^{-1}[X]$$

for every $X \in \mathcal{P}(S)$. We note that $\mathcal{F}^+$ is a normal modal Boolean algebra.

Complex algebras are the concrete modal Boolean algebras that algebraize relational semantics [25, Theorem 5.25]. By means of complex algebras one can construct a modal Boolean algebra from a given Kripke frame. For the other direction we need to construct the *ultrafilter frame* [25, Definition 5.34]. By transforming this frame in a complex algebra we get the *Jònson-Tarski Theorem* underlying the algebraization of modal logic:

> *Every modal Boolean algebra can be embedded in the complex algebra of its ultrafilter frame [25, Theorem 5.43].*

Convention 3.3.5 *Throughout this thesis, when we say complex algebra of a Kripke model $\mathcal{M} = (\mathcal{F}, V)$, we mean complex algebra of the underlying Kripke frame $\mathcal{F}$.* ⊣

As a final step to show that the class of complex algebras algebraizes the semantics of modal logic, we need to interpret modal formulas in modal Boolean algebras. To this end, we introduce *algebraic models*.

Definition 3.3.6 (Algebraic model) *An algebraic model is a pair $\mathcal{A} = \langle \mathbb{A}, V \rangle$ where $\mathbb{A}$ is a modal Boolean algebra and $V : \mathsf{At} \to \mathbb{A}$ is a valuation that assigns an element from $\mathbb{A}$ to each atomic proposition.*

Also, given a Kripke model $\mathcal{M} = \langle \mathcal{F}, V \rangle$ with $V : \mathsf{At} \to \mathcal{P}(S)$, the algebraic model associated with $\mathcal{M}$ is the tuple $\mathcal{A} = \langle \mathbb{A}, V \rangle$ where $\mathbb{A}$ is the complex algebra of $\mathcal{F}$. ⊣

We will rely on the duality between Kripke frames and normal modal Boolean algebras to define the algebraic semantics of action model logic and refinement action modal logic.

3.4 Epistemic updates on algebras

In this section, we will report on how to define epistemic updates on modal Boolean algebras (Subsection 3.4.1), and then, we will recall the algebraic semantics of action model logic (Subsection 3.4.3). Our presentation is a summary of the method presented in [105, 114].

3.4.1 Methodology

We first describe a two-step account of the product update construction on Kripke models from Def. 3.2.9, and then the mathematical steps to compute the updated algebra from [105].

Throughout the present subsection, we fix a Kripke model $\mathcal{M} = \langle S, R, V \rangle$ and an epistemic action α_k where $\alpha = (K, R_\alpha, Pre_\alpha)$. The product update $\mathcal{M}_\alpha$ defined in Section 3.2 can be built in two steps as follows.

<u>STEP 1</u> We define the following intermediate model

$$\coprod_\alpha \mathcal{M} = \langle \coprod_K S, R \times R_\alpha, \coprod_\alpha V \rangle$$

where

- $\coprod_K S \simeq S \times K$ is the $|K|$-fold coproduct S, which is set-isomorphic to the Cartesian product $S \times K$,
- $R \times R_\alpha$ is the binary relation on $\coprod_K S$ defined as

$$(s', k') \in R \times R_\alpha((s', k')) \quad \text{iff} \quad s' \in R(s') \;\text{ and }\; k' \in R_\alpha(k),$$

- $\coprod_\alpha V : \mathsf{At} \to \mathcal{P}(\coprod_K S)$ such that for every $p \in \mathsf{At}$

$$\coprod_\alpha V(p) = \coprod_\alpha (V(p)) = V(p) \times K.$$

<u>Step 2.</u> $\mathcal{M}_\alpha$ is the submodel of $\coprod_\alpha \mathcal{M}$ that contains exactly all the tuples $(s,k) \in \coprod_K S$ such that $(\mathcal{M}, s) \models Pre_\alpha(k)$.

This two-step-account of the product update construction can be seen as a pseudo-coproduct followed by taking a submodel as illustrated by the following diagram

$$\mathcal{M} \hookrightarrow \coprod_\alpha \mathcal{M} \hookleftarrow \mathcal{M}_\alpha.$$

This perspective makes it possible to use the duality between products and coproducts in category theory (cf. [44, 8]): coproducts can be dually characterised as products, and subobjects as quotients. Using this result, the update of $\mathcal{M}$ with the action model α, regarded as a "subobject after coproduct" concatenation, can be dually characterised on its algebraic counterpart $\langle \mathbb{A}, V\rangle$ by means of a "quotient after product" concatenation, as illustrated in the following diagram:

$$\mathbb{A} \twoheadleftarrow \prod_\alpha \mathbb{A} \twoheadrightarrow \mathbb{A}^\alpha. \tag{3.3}$$

Indeed, the pseudo-coproduct $\coprod_\alpha \mathcal{M}$ is dually characterised as a *pseudo-product* $\prod_\alpha \mathbb{A}$ and an appropriate *quotient* of $\prod_\alpha \mathbb{A}$ is then taken to dually characterise the submodel step. This construction we now define.

3.4.2 Dual characterisation of the intermediate model

Definition 3.4.1 (Action model on modal Boolean algebras) *Given a modal Boolean algebra $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond, \mathbf{0}, \mathbf{1}\rangle$, we define an* action model *over $\mathbb{A}$ as a tuple $a = (K, R_a, Pre_a)$ such that K is a finite nonempty set, $R_\alpha \subseteq K \times K$ and $Pre_a : K \to \mathbb{A}$. As for Kripke models, one can define pointed action models (a, k) over $\mathbb{A}$ with $k \in a$ denoted a_k. An action model over a Boolean algebra $\mathbf{A} = \langle A, \wedge, \vee, \neg, \mathbf{0}, \mathbf{1}\rangle$ is defined in a similar way except that Pre_a maps every $k \in K$ to an element from $\mathbf{A}$.*

Given Kripke model $\mathcal{M} = (\mathcal{F}, V)$, an action model $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle$ induces a corresponding action model $a = \langle K, R_a, Pre_a\rangle$ over the complex algebra $\mathbb{A}$ of $\mathcal{F}$, where $R_a = R_\alpha$ and $Pre_a : K \to \mathbb{A}$ is defined as $Pre_a(k) = [\![Pre_\alpha(k)]\!]_\mathcal{M}$ where $[\![-]\!]_\mathcal{M} : \mathcal{L}_{\Box\alpha} \to \mathbb{A}$ is an extension map associated with $\mathcal{M}$. ⊣

Definition 3.4.2 (Intermediate algebra) *For every modal Boolean algebra $\mathbb{A} = \langle \mathbf{A}, \Diamond\rangle$ and every action model $a = \langle K, R_a, Pre_a\rangle$ over $\mathbb{A}$, we define the* intermediate modal Boolean algebra *$\prod_a \mathbb{A} = \langle \mathbf{A}^K, \Diamond_{\prod_a \mathbf{A}}\rangle$ as a modal expansion of the $|K|$-fold product $\mathbf{A}^K$ of $\mathbf{A}$, which is the Boolean algebra that has as its carrier the set A^K of set maps $f : K \to A$, and in which the Boolean operations are defined pointwise, i.e., for all $f, g : K \to A$:*

$$\begin{aligned}(f \vee_{\mathbf{A}^K} g)(k) &:= f(k) \vee_\mathbb{A} g(k),\\ (f \wedge_{\mathbf{A}^K} g)(k) &:= f(k) \wedge_\mathbb{A} g(k),\\ (\neg_{\mathbf{A}^K} f)(k) &:= \neg_\mathbb{A} f(k),\\ c_{\mathbf{A}^K}(k) &:= c_\mathbb{A}. \qquad (c \in \{\mathbf{0}, \mathbf{1}\})\end{aligned}$$

The modalities $\Diamond_{\prod_a \mathbb{A}} f : K \to A$ *and* $\Box_{\prod_a \mathbb{A}} f : K \to A$ *are defined by:*

$$(\Diamond_{\prod_a \mathbb{A}} f)(k) = \bigvee \{\Diamond_{\mathbb{A}} f(k') \ : \ k' \in R_a(k)\}, \tag{3.4}$$

$$(\Box_{\prod_a \mathbb{A}} f)(k) = \bigwedge \{\Box_{\mathbb{A}} f(k') \ : \ k' \in R_a(k)\}. \tag{3.5}$$

We refer to [105, Section 3.1] for an extensive justification of the above definition. Similar definitions apply to $\blacklozenge_{\prod_a \mathbb{A}}$ and $\blacksquare_{\prod_a \mathbb{A}}$.

Any property of $\Diamond_{\mathbb{A}}$ and $\Box_{\mathbb{A}}$ will be inherited by $\Diamond_{\prod_a \mathbb{A}}$ and $\Box_{\prod_a \mathbb{A}}$, respectively. Hence, $\Diamond_{\prod_a \mathbb{A}}$ and $\Box_{\prod_a \mathbb{A}}$ are normal and $\Box_{\prod_a \mathbb{A}} = \neg \Diamond_{\prod_a \mathbb{A}} \neg$ (cf. [105]). Similarly, it can be shown that $\prod_a \mathbb{A} = \langle \mathbf{A}^K, \Diamond_{\prod_a \mathbb{A}}, \blacksquare_{\prod_a \mathbb{A}} \rangle$ is a modal Boolean algebra, the so-called *tense intermediate algebra*.

Moreover, in the case that $\mathbb{A} = \mathcal{F}^+$ is the complex algebra of a given Kripke model $\mathcal{M} = (\mathcal{F}, V)$ and a is an action model over $\mathbb{A}$ induced by an action model α over $\mathcal{L}_{\Box\alpha}$, the intermediate algebra $\prod_a \mathbb{A}$ is isomorphic to the complex algebra of the intermediate model $\coprod_a \mathcal{M}$ [105, Proposition 3.1].

At this point, we dually characterise the product update by defining a quotient structure over the intermediate algebra.

(Pseudo-)quotient of the intermediate algebra. Let $\mathbb{A}$ be a modal Boolean algebra, and let $a = \langle K, R_a, Pre_a \rangle$ be an action model over $\mathbb{A}$. The equivalence relation $\equiv_a$ on $\prod_a \mathbb{A} = \langle \mathbf{A}^K, \Diamond_{\prod_a \mathbb{A}} \rangle$ is defined as follows: for all $f, g \in \mathbf{A}^K$,

$$f \equiv_a g \ \text{ iff } \ f \wedge Pre_a = g \wedge Pre_a. \tag{3.6}$$

The equivalence relation $\equiv_a$ is a congruence with respect to Boolean operations. We denote by $\mathbf{A}_a$ the quotient Boolean algebra $\mathbf{A}^K/{\equiv_a}$, and by $[f]_a$ the equivalence class of every $f \in \mathbf{A}^K$. The subscript will be dropped whenever it causes no confusion. However, $\equiv_a$ is *not* compatible with the modalities, indeed $f \equiv_a g$ does not imply that $\Diamond f \equiv_a \Diamond g$. Hence, we need to choose a definition for the modalities on $\mathbf{A}_a$. To this end, $\Diamond^a$ and $\Box^a$, for every $f \in \mathbf{A}^K$, are defined as follows.

$$\Diamond^a [f] := [\Diamond_{\prod_a \mathbb{A}} (f \wedge Pre_a)] \tag{3.7}$$

$$\Box^a [f] := [\Box_{\prod_a \mathbb{A}} (f \to Pre_a)]. \tag{3.8}$$

The modalities $\Diamond^a$ and $\Box^a$ are normal and $\Box^a = \neg \Diamond^a \neg$. So $\mathbb{A}_a = \langle \mathbf{A}_a, \Diamond^a \rangle$ belongs to **MBA**. Moreover, $\mathbb{A}_a$ behaves in the desired way, in the sense that if $\mathbb{A} = \mathcal{F}^+$ for some Kripke frame $\mathcal{F}$, we get that $\mathbb{A}_a \cong (\mathcal{F}_\alpha)^+$, where $\mathcal{F}_\alpha$ is the underlying Kripke frame of the product update $\mathcal{M}_\alpha$. The same definition as (3.7) and (3.8) applies to $\blacklozenge^a$ and $\blacksquare^a$, respectively.

Definition 3.4.3 (Updated algebra) *Let* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Diamond, \mathbf{0}, \mathbf{1} \rangle$ *be a modal Boolean algebra, and let a be an action model over* $\mathbb{A}$*. The pseudo-quotient modal Boolean algebra* $\mathbb{A}_a = \langle \mathbf{A}_a, \Diamond^a \rangle$ *defined as above, is called the* updated algebra*. Similarly, we define the* tense updated algebra *by expanding* $\mathbb{A}_a$ *with* $\blacklozenge^a$. ⊣

3.4.3 Algebraic semantics of action model logic

In this subsection, we briefly recall the algebraic semantics of action model logic proposed by [105, 114].

In the previous subsection we introduced the algebraic counterparts of intermediate model and product update model. We now turn these algebras into models.

Definition 3.4.4 (Intermediate algebraic model) *Let $\mathcal{A} = \langle \mathbb{A}, V \rangle$ be an algebraic model (Def. 3.3.6), and let $\alpha = \langle K, R_\alpha, Pre_\alpha \rangle$ be an action model. Let a be the action model induced by α via $[\![-]\!]_{\mathcal{M}}$. The* intermediate algebraic model $\prod_\alpha \mathcal{A}$ *is defined as*

$$\prod_\alpha \mathcal{A} := \langle \prod_a \mathbb{A}, \prod_a V \rangle$$

where, (i) $\prod_a \mathbb{A}$ is the intermediate modal Boolean algebra, and (ii) for any $p \in \mathsf{At}$, the valuation $(\prod_a V)(p) : K \to \prod_a \mathbb{A}$ is defined by taking $((\prod_a V)(p))(k) = V(p)$. ⊣

Let us look at Diagram (3.3) again. Assume $\mathbb{A}$ is a modal Boolean algebra and $a = (K, R_a, Pre_a)$ is an action model over $\mathbb{A}$. As $\mathbb{A}_a$ is the pseudo-quotient of $\prod_a \mathbb{A}$, there is a *quotient map* $q : \prod_a \mathbb{A} \to \mathbb{A}_a$ mapping each $f \in \prod_a \mathbb{A}$ to the equivalence class $[f]_a$ in $\mathbb{A}_a$. It is easy to check that the quotient map is monotone. On the other hand, one can define a mapping $i' : \mathbb{A}_a \to \prod_a \mathbb{A}$ by taking $i'([f]) = f \wedge Pre_a$ for all $[f] \in \mathbb{A}_a$, which is monotone, as well. It is not difficult to check that the pair (q, i') forms an adjunction. In addition, for each point k of K, the projection on the k-indexed coordinate $\pi_k : \prod_a \mathbb{A} \to \mathbb{A}$ maps every $f \in \prod_a \mathbb{A}$ to $f(k)$. Let us summarize what we have learned so far in the following diagram.

$$\mathbb{A} \overset{}{\underset{\pi_k}{\twoheadleftarrow}} \prod_a \mathbb{A} \underset{q}{\overset{i'}{\rightleftarrows}} \mathbb{A}_a \tag{3.9}$$

Definition 3.4.5 (Updated algebraic model) *Let $\mathcal{A} = \langle \mathbb{A}, V \rangle$ be an algebraic model, $\alpha = \langle K, R_\alpha, Pre_\alpha \rangle$ is an action model. Let a be the action model induced by α via V. The* updated algebraic model $\mathcal{A}_\alpha$ *is defined as*

$$\mathcal{A}_\alpha := \langle \mathbb{A}_a, V_a \rangle$$

where $V_a : \mathsf{At} \to \mathbb{A}_a$ is the map such that $V_a(p) = q \circ \prod_a V(p) = [\prod_a V(p)]_a$. ⊣

Finally, here is how we obtain an algebraic semantics for **AML**.

Definition 3.4.6 (Algebraic semantics of AML) *For every algebraic model $\mathcal{A} = \langle \mathbb{A}, V \rangle$, the extension map $[\![-]\!]_{\mathcal{A}} : \mathcal{L}_{\Box\alpha} \to \mathbb{A}$ is defined recursively as follows:*

$$\begin{aligned}
[\![p]\!]_{\mathcal{A}} &:= V(p) \\
[\![\neg\varphi]\!]_{\mathcal{A}} &:= \neg_{\mathbb{A}} [\![\varphi]\!]_{\mathcal{A}} \\
[\![\Box\varphi]\!]_{\mathcal{A}} &:= \Box_{\mathbb{A}} [\![\varphi]\!]_{\mathcal{A}} \\
[\![\varphi \bullet \psi]\!]_{\mathcal{A}} &:= [\![\varphi]\!]_{\mathcal{A}} \bullet_{\mathbb{A}} [\![\psi]\!]_{\mathcal{A}} \qquad (\text{ for } \bullet \in \{\vee, \wedge, \to\}) \\
[\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{A}} &:= [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \wedge_{\mathbb{A}} \pi_k \circ i'([\![\varphi]\!]_{\mathcal{A}^\alpha}) \\
[\![[\alpha_k]\varphi]\!]_{\mathcal{A}} &:= [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \to_{\mathbb{A}} \pi_k \circ i'([\![\varphi]\!]_{\mathcal{A}^\alpha})
\end{aligned}$$

where $\alpha = \langle K, R_a, Pre_a \rangle \in \mathcal{AM}_{AML}$ is an action model and $k \in K$. ⊣

Lemma 3.4.7 *[105, Lemma 7.5,7.6,7.8] Let $\mathcal{A} = (\mathbb{A}, V)$ be an algebraic model and let α be a pointed action model.*

1. $[\![\langle\alpha_k\rangle(\varphi \vee \psi)]\!]_{\mathcal{A}} = [\![\langle\alpha_k\rangle\varphi \vee \langle\alpha_k\rangle\varphi]\!]_{\mathcal{A}}$.
2. $[\![[\alpha_k](\varphi \vee \psi)]\!]_{\mathcal{A}} = [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \to ([\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{A}} \vee [\![\langle\alpha_k\rangle\psi]\!]_{\mathcal{A}})$.

3. $[\![\langle\alpha_k\rangle(\varphi \wedge \psi)]\!]_{\mathcal{A}} = [\![\langle\alpha_k\rangle\varphi \wedge \langle\alpha_k\rangle\varphi]\!]_{\mathcal{A}}$.

4. $[\![[\alpha_k](\varphi \wedge \psi)]\!]_{\mathcal{A}} = [\![[\alpha_k]\varphi \wedge [\alpha_k]\psi]\!]_{\mathcal{A}}$.

5. $[\![\langle\alpha_k\rangle\Box\varphi]\!]_{\mathcal{A}} = [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \wedge \bigwedge\{\Box[\![[\alpha_j]\varphi]\!]_{\mathcal{A}} : j \in R_\alpha(k)\}$.

6. $[\![\langle\alpha_k\rangle\Diamond\varphi]\!]_{\mathcal{A}} = [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \wedge \bigvee\{\Diamond[\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{A}} : j \in R_\alpha(k)\}$.

7. $[\![\langle\alpha_k\rangle\neg\varphi]\!]_{\mathcal{A}} = [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \wedge \neg[\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{A}}$.

8. $[\![[\alpha_k]\neg\varphi]\!]_{\mathcal{A}} = [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \rightarrow \neg[\![[\alpha_k]\varphi]\!]_{\mathcal{A}}$. ⊣

3.5 Algebraic semantics of refinement action modal logic

In this section, we present our main result, namely an algebraic semantics for **RAML**. It is an extension of the algebraic semantics for **AML** of the previous section. First we introduce the notion of *refinement morphism*. It is the analogue on normal boolean algebras with operators of the notion of refinement between Kripke models. Then we define the notion of *refinement algebra*. This is used in the definition of the algebraic semantics of refinement action modal logic. Finally we prove that $\mathbb{RAML}$ is sound and complete with respect to this semantics.

In the sequel, we assume $\mathcal{A} = \langle\mathbb{A}, V\rangle$ is an algebraic model in which $\mathbb{A}$ is a complete modal Boolean algebra.

Let us expand on how the refinement relation between Kripke models can be dualized as a relation between modal Boolean algebras. Let $\mathfrak{R}$ be a non-empty binary relation between $\mathcal{M} = \langle S, R, V\rangle$ and $\mathcal{M}' = \langle S', R', V'\rangle$. Let $s \in S$, then the **(Zag)** condition from Definition 3.2.6 can be rewritten as follows:

$$R'[\mathfrak{R}(s)] \subseteq \mathfrak{R}\,[R(s)]. \tag{3.10}$$

In light of the Jónsson-Tarski theorem, (3.10) can be dually characterised on algebras by mean of what we call *refinement morphisms*. We prove that refinement morphisms are right adjoints.

Definition 3.5.1 (Refinement morphism) *Let $\mathbb{A}$ and $\mathbb{A}'$ be two modal Boolean algebras. A map $g : \mathbb{A} \to \mathbb{A}'$ is a* refinement morphism *if (i) it preserves* $\mathbf{0}$ *and* $\vee$*, i.e.,* $g(\mathbf{0}_{\mathbb{A}}) = \mathbf{0}_{\mathbb{A}'}$ *and* $g(x \vee_{\mathbb{A}} y) = g(x) \vee_{\mathbb{A}'} g(y)$ *for every* $x, y \in \mathbb{A}$*, and (iii) satisfies the following inequality*

$$(\textbf{\textit{AlgZag}}) \qquad \blacklozenge_{\mathbb{A}'} \circ g \;\leq\; g \circ \blacklozenge_{\mathbb{A}}$$

where $\blacklozenge_{\mathbb{A}}$ *and* $\blacklozenge_{\mathbb{A}'}$ *are the left adjoints of* $\Box_{\mathbb{A}}$ *and* $\Box_{\mathbb{A}'}$*, respectively.* ⊣

The **AlgZag** inequality is the dual of the **Zag** condition in the refinement relation (cf. page 40).

Recall that by Theorem 3.2.19, for every formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ and every Kripke model $\mathcal{M}$, there is a multi-pointed action model $\alpha^\varphi = (\langle K^\varphi, R^\varphi_\alpha, Pre^\varphi_\alpha\rangle, K^\varphi)$ such that $\mathcal{M}_{\alpha^\varphi}$ is a refinement of $\mathcal{M}$ and $(\mathcal{M}, s) \models \exists\varphi$ iff $(\mathcal{M}_{\alpha^\varphi}, (s, k)) \models \varphi$ for some $k \in K^\varphi$. By the dual characterisation introduced in Subsection 3.4.1, for a given formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, and for all algebraic model $\mathcal{A} = \langle\mathbb{A}, V\rangle$, $\alpha^\varphi = \langle K^\varphi, R^\varphi_\alpha, Pre^\varphi_\alpha\rangle$ induces an action model $a^\varphi = \langle K^\varphi, R^\varphi_a, Pre^\varphi_a\rangle$ over $\mathbb{A}$, and we obtain the updated algebra $\mathbb{A}^\varphi := \mathbb{A}_{a^\varphi}$ (see Definition 3.4.3), and a pair of maps (g^φ, f^φ) as follows:

$$
\begin{aligned}
g^{\varphi} &: \mathbb{A} \to \mathbb{A}^{\varphi} & f^{\varphi} &: \mathbb{A}^{\varphi} \to \mathbb{A} \\
x &\mapsto [h_x] & [h] &\mapsto \bigvee_{k \in K^{\varphi}} (h(k) \wedge Pre^{\varphi}_a(k))
\end{aligned}
\tag{3.11}
$$

where $h_x : K^{\varphi} \to \mathbb{A}$ is the map defined by $h_x(k) := x \wedge Pre^{\varphi}_a(k)$ for every $k \in K^{\varphi}$. We denote by $\mathcal{A}^{\varphi}$ the updated algebraic model $\mathcal{A}_{\alpha^{\varphi}} = (\mathbb{A}^{\varphi}, V^{\varphi})$.

Remark 3.5.2 *One can easily show that* $f^{\varphi}([h]) = \bigvee_{k \in K^{\varphi}} \pi^{\varphi}_k \circ i'^{\varphi}([h])$ *for every* $[h] \in \mathbb{A}^{\varphi}$, *where* $\pi^{\varphi}_k : \prod_{a^{\varphi}} \mathbb{A} \to \mathbb{A}$ *and* $i'^{\varphi} : \mathbb{A}^{\varphi} \to \prod_{a^{\varphi}} \mathbb{A}$ *are defined in the diagram* (3.9). ⊣

Lemma 3.5.3 *For all algebraic model* $\mathcal{A} = \langle \mathbb{A}, V \rangle$ *and all formula* $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, *we have that*

1. *the map* g^{φ} *is a refinement morphism,*
2. *the map* f^{φ} *preserves arbitrary joins,*
3. $f^{\varphi} \dashv g^{\varphi}$.
4. $f^{\varphi}([\![\psi]\!]_{\mathcal{A}^{\varphi}}) = [\![\langle \alpha^{\varphi}_{K^{\varphi}} \rangle \psi]\!]_{\mathcal{A}}$, *for all* $\varphi, \psi \in \mathcal{L}_{\Box\alpha\forall}$, *where* $\alpha^{\varphi} = (K^{\varphi}, R^{\varphi}_{\alpha}, Pre^{\varphi}_{\alpha})$ *is an action model associated with* φ *and* $\mathcal{A}^{\varphi} = (\mathbb{A}^{\varphi}, V^{\varphi})$ *is the update algebra.* ⊣

Proof *Item 1.* It follows from [105, Fact 3.7(4)] and Lemma 3.2.20 that for every $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, g^{φ} preserves $\mathbf{0}$ and $\vee$. It only remains to check that g^{φ} satisfies the (**AlgZag**) condition. To this end, we first note that for every modal Boolean algebra $\mathbb{A}$, and every action model $a = \langle K, R_a, Pre_a \rangle$ over $\mathbb{A}$, [105, Fact 3.5(2)] shows that for every $h, h' \in \prod_a \mathbb{A}$,

$$
[h]_a \leq [h']_a \iff h \wedge Pre_a \leq h' \wedge Pre_a \tag{3.12}
$$

Now, let $x \in \mathbb{A}$. It follows from the definition of g^{φ} and (3.7) that

$$
(\blacklozenge_{\mathbb{A}^{\varphi}} \circ g^{\varphi})(x) = \left[\blacklozenge_{\prod_{\alpha} \mathbb{A}}(h_x \wedge Pre^{\varphi}_a)\right]_a \qquad \textit{and} \qquad g^{\varphi} \circ \blacklozenge_{\mathbb{A}}(x) = [\blacklozenge_{\mathbb{A}} x]_a .
$$

We will show that for every $k \in K^{\varphi}$,

$$
\left(\blacklozenge_{\prod_{\alpha} \mathbb{A}}(h_x \wedge Pre^{\varphi}_a) \wedge Pre^{\varphi}_a\right)(k) \leq (\blacklozenge_{\mathbb{A}} x \wedge Pre^{\varphi}_a)(k).
$$

Assume $k \in K^{\varphi}$, we have

$$
\begin{aligned}
&\left(\blacklozenge_{\prod_{\alpha} \mathbb{A}}(h_x \wedge Pre^{\varphi}_a) \wedge Pre^{\varphi}_a\right)(k) \\
&= \bigvee \{\blacklozenge_{\mathbb{A}}(h_x \wedge Pre^{\varphi}_a)(j) \;:\; j \in R_a(k)\} \wedge Pre^{\varphi}_a(k) && (3.4) \\
&\leq \bigvee \{\blacklozenge_{\mathbb{A}} h_x(j) \wedge \blacklozenge_{\mathbb{A}} Pre^{\varphi}_a(j) \;:\; j \in R_a(k)\} \wedge Pre^{\varphi}_a(k) && (\blacklozenge(x \wedge y) \leq \blacklozenge x \wedge \blacklozenge y) \\
&= \bigvee \{\blacklozenge_{\mathbb{A}}(x \wedge Pre^{\varphi}_a(j)) \wedge \blacklozenge_{\mathbb{A}} Pre^{\varphi}_a(j) \;:\; j \in R_a(k)\} \wedge Pre^{\varphi}_a(k) && (\text{Def. } h_x) \\
&= \bigvee \{\blacklozenge_{\mathbb{A}} x \wedge \blacklozenge Pre^{\varphi}_a(j) \wedge \blacklozenge_{\mathbb{A}} Pre^{\varphi}_a(j) \;:\; j \in R_a(k)\} \wedge Pre^{\varphi}_a(k) && (\blacklozenge(x \wedge y) \leq \blacklozenge x \wedge \blacklozenge y) \\
&\leq \bigvee \{\blacklozenge_{\mathbb{A}} x \wedge \blacklozenge_{\mathbb{A}} Pre^{\varphi}_a(j) \;:\; j \in R_a(k)\} \wedge Pre^{\varphi}_a(k) && (x \wedge x = x) \\
&\leq \bigvee \{\blacklozenge_{\mathbb{A}} x \;:\; j \in R_a(k)\} \wedge Pre^{\varphi}_a(k) && (\blacklozenge x \wedge \blacklozenge y \leq \blacklozenge x) \\
&\leq \blacklozenge_{\mathbb{A}} x \wedge Pre^{\varphi}_a(k) && (K^{\varphi} \text{ finite, } \underbrace{x \vee x \vee ... \vee x}_{n-times} = x)
\end{aligned}
$$

Therefore, g^φ is a refinement morphism.

Item 2. We first show that f^φ is monotone. Let $[h]_a, [h'] \in \mathbb{A}^\varphi$ be such that $[h] \leq [h']$. It follows from (3.12) that $h(k) \wedge Pre^\varphi_a(k) \leq h'(k) \wedge Pre^\varphi_a(k)$ for every $k \in K^\varphi$. Hence, $\bigvee_{k\in K}(h(k) \wedge Pre^\varphi_a(k)) \leq \bigvee_{k\in K}(k(k) \wedge Pre^\varphi_a(k))$, which shows that $f^\varphi([h]) \leq f^\varphi([k])$. Thus, f^φ preserves $\mathbf{0}$. In order to prove that f^φ preserves arbitrary joins, we have to show that $f^\varphi(\bigvee_{i\in I}[h_i]) = \bigvee_{i\in I} f^\varphi([h_i])$, where $[h_i] \in \mathbb{A}^\varphi$ and i is taken from some index set I. First, as $\mathbb{A}$ is complete modal Boolean algebra, the intermediate and updated algebras are complete, as well. Thus, $\bigvee_{i\in I}[h_i] = [\bigvee_{i\in I} h_i]$, where $[h_i] \in \mathbb{A}^\varphi$. So, we have

$$
\begin{aligned}
f^\varphi\left(\bigvee_{i\in I}[h_i]\right) &= f^\varphi\left(\left[\bigvee_{i\in I} h_i\right]\right) \\
&= \bigvee_{k\in K^\varphi}\left((\bigvee_{i\in I} h_i(k)) \wedge Pre^\varphi_a(k)\right) && \text{(Def. of } f^\varphi) \\
&= \bigvee_{k\in K^\varphi}\bigvee_{i\in I}(h_i(k) \wedge Pre^\varphi_a(k)) && (\textstyle\bigvee_{i\in I} x_i \wedge y = \bigvee_{i\in I}(x_i \wedge y)) \\
&= \bigvee_{i\in I}\bigvee_{k\in K^\varphi}(h_i(k) \wedge Pre^\varphi_a(k)) && (K^\varphi \text{ finite}) \\
&= \bigvee_{i\in I}(f^\varphi[h_i])\,. && \text{(Def. } f^\varphi)
\end{aligned}
$$

For *item 3* we need to show that for every $x \in \mathbb{A}$, and $[h] \in \mathbb{A}^\varphi$, $[h] \leq g^\varphi(x)$ iff $f^\varphi([h]) \leq x$. By (3.12) and the definition of g^φ and f^φ, it is equivalent to show that $h \wedge Pre^\varphi_a \leq h_x \wedge Pre^\varphi_a$ iff $\bigvee_{k\in K^\varphi} h(k) \wedge Pre^\varphi_a(k) \leq x$. Let $x \in \mathbb{A}$, then for all $k \in K^\varphi$ we have

$$
\begin{aligned}
&h(k) \wedge Pre^\varphi_a(k) \leq h_x(k) \wedge Pre^\varphi_a(k) \\
&\qquad \text{iff} \quad h(k) \wedge Pre^\varphi_a(k) \leq (x \wedge Pre^\varphi_a(k)) \wedge Pre^\varphi_a(k) && \text{(Def. } h_x) \\
&\qquad \text{iff} \quad h(k) \wedge Pre^\varphi_a(k) \leq x \wedge Pre^\varphi_a(k) \leq x \\
&\qquad \text{iff} \quad \bigvee_{k\in K^\varphi}(h \wedge Pre^\varphi_a)(k) \leq x && \text{(Def. } \textstyle\bigvee) \\
&\qquad \text{iff} \quad f^\varphi([h]) \leq x. && \text{(Def. } f^\varphi)
\end{aligned}
$$

Item 4. Without loss of generality, from Fact 3.2.20 we may assume $\varphi \in \mathcal{L}_{\Box\alpha}$. By Def. 3.4.6, we have $[\![\langle\alpha^\varphi_{K^\varphi}\rangle\psi]\!]_{\mathcal{A}} = \bigvee_{k\in K^\varphi}[\![Pre^\varphi_\alpha(k)]\!]_{\mathcal{A}} \wedge \pi^\varphi_k \circ i'^\varphi([\![\psi]\!]_{\mathcal{A}^\varphi})$. For the sake of a smooth presentation, let $[h^\psi] := [\![\psi]\!]_{\mathcal{A}^\varphi}$. As noted in Remark 3.5.2, $\pi^\varphi_k \circ i'^\varphi([h^\psi]) = Pre^\varphi_a(k) \wedge h^\psi(k)$. Since $[\![Pre^\varphi_\alpha(k)]\!]_{\mathcal{A}} = Pre^\varphi_a(k)$, we have $[\![\langle\alpha^\varphi_{K^\varphi}\rangle\psi]\!]_{\mathcal{A}} = \bigvee_{k\in K^\varphi} Pre^\varphi_a(k) \wedge h^\psi(k)$. Now, we compute $f^\varphi([\![\psi]\!]_{\mathcal{A}^\varphi})$. By definition of f^φ, we have

$$
f^\varphi([\![\psi]\!]_{\mathcal{A}^\varphi}) = \bigvee_{k\in K^\varphi} h^\psi(k) \wedge Pre^\varphi_a(k). \tag{3.13}
$$

where $[h^\psi] := [\![\psi]\!]_{\mathcal{A}^\varphi}$. Hence $f^\varphi([\![\psi]\!]_{\mathcal{A}^\varphi}) = [\![\langle\alpha^\varphi_{K^\varphi}\rangle\psi]\!]_{\mathcal{A}}$, as desired. □

The next lemma illustrates that the refinement relation between two Kripke models induces a refinement morphism between the complex algebras associated with them, and vice versa.

Lemma 3.5.4 *Let $\mathcal{M} = \langle S, R, V\rangle$ and $\mathcal{M}' = (S', R', V')$ be two Kripke models with underlying frames $\mathcal{F}$ and $\mathcal{F}'$, respectively. Assume that $\mathcal{A} = \langle\mathbb{A}, V\rangle$ and $\mathcal{A}' = \langle\mathbb{A}', V'\rangle$ are algebraic models* *such that* $\mathbb{A} = \mathcal{F}^+$ *and* $\mathbb{A}' = \mathcal{F}'^+$. *Then,* $(\mathcal{M}, s) \succeq (\mathcal{M}', s')$ *iff there exists a refinement morphism* $g : \mathbb{A} \to \mathbb{A}'$ *with* $\{s'\} \leq g(\{s\})$ *such that* $g(V(p)) \leq V'(p)$ *and* $g^{-1}(V'(p)) \leq' V(p)$ *for every* $p \in \mathsf{At}$. ⊣

Proof For the direction from left to right, assume that $(\mathcal{M}, s) \succeq (\mathcal{M}', s')$. There is a refinement $\mathfrak{R}$ from $\mathcal{M}$ to $\mathcal{M}'$ such that $s\mathfrak{R}s'$. We define the map $g : \mathbb{A} \to \mathbb{A}'$ by $g(X) = \mathfrak{R}[X]$, for every $X \subseteq S$. It is easy to see that g is a refinement morphism with $\{s'\} \leq g(\{s\})$. It remains to show that for every $p \in \mathsf{At}$, $g(V(p)) \leq V'(p)$ and $g^{-1}(V'(p)) \leq' V(p)$. Let $p \in \mathsf{At}$. We first note that $\leq$ and $\leq'$ over $\mathbb{A}$ and $\mathbb{A}'$ are interpreted as inclusion over $\mathcal{P}(S)$ and $\mathcal{P}(S')$, respectively. So, if $s' \in g(V(p))$, then by the definition of g, $s' \in \mathfrak{R}[V(p)]$. It implies that there is $s \in V(p)$ such that $s' \in \mathfrak{R}(s)$. As $\mathfrak{R}$ is a refinement, by (**Atoms**), $s \in V(p)$ iff $s' \in V'(p)$. Thus, $s' \in V'(p)$. The other inequality is proved in a similar way.

For the other direction, let $g : \mathbb{A} \to \mathbb{A}'$ be a refinement morphism as in the statement of the lemma. Define a binary relation on $S \times S'$ by: $\mathfrak{R} = \{(s, s') \in S \times S' \ : \ s' \in f(\{s\})\}$. It is easy to see that $\mathfrak{R}$ is a refinement. □

3.5.1 Algebraic model of refinement modality

We aim at proposing an algebraic semantics for the refinement quantifiers $\exists$, i.e. for all algebraic model $\mathcal{A} = \langle \mathbb{A}, V \rangle$, we want to find a modal Boolean algebra $\mathfrak{A}_{\mathbb{A}}$ and a map $F : \mathfrak{A}_{\mathbb{A}} \to \mathbb{A}$ such that for any $\varphi \in \mathcal{L}_{\Box\alpha\forall}$,

$$[\![\exists\varphi]\!]_{\mathcal{A}} = F([\![\varphi]\!]_{\mathfrak{A}_{\mathbb{A}}}).$$

To do so, we introduce a modal Boolean algebra $\mathfrak{A}_{\mathbb{A}}$ such that for each $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, $\mathbb{A}^{\varphi}$ is a subalgebra of $\mathfrak{A}_{\mathbb{A}}$.

Definition 3.5.5 (Refinement algebra) *Let* $\mathbb{A}$ *be a modal Boolean algebra, and let* $(\mathbb{A}^{\varphi})_{\varphi \in \mathcal{L}_{\Box\alpha\forall}}$ *be the family of all updated algebras. The* refinement algebra *of* $\mathbb{A}$ *is defined as the product of this family:*

$$\mathfrak{A}_{\mathbb{A}} := \prod_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} \mathbb{A}^{\varphi}.$$

The elements of $\mathfrak{A}_{\mathbb{A}}$ *are tuples* $(x^{\varphi})_{\varphi \in \mathcal{L}_{\Box\alpha\forall}}$ *where* $x^{\varphi} \in \mathbb{A}^{\varphi}$. *When there is no risk of confusion, we write* $(x^{\varphi})_{\varphi}$ *instead of* $(x^{\varphi})_{\varphi \in \mathcal{L}_{\Box\alpha\forall}}$ *and* $\mathfrak{A}$ *instead of* $\mathfrak{A}_{\mathbb{A}}$. *The operations are defined coordinatewise as follows: for all* $(x^{\varphi})_{\varphi}, (y^{\varphi})_{\varphi} \in \mathfrak{A}$,

Constants	**Negation**
$\mathbf{0}_{\mathfrak{A}} = (\mathbf{0}^{\varphi})_{\varphi}$, $\mathbf{1}_{\mathfrak{A}} = (\mathbf{1}^{\varphi})_{\varphi}$	$\neg_{\mathfrak{A}}(x^{\varphi})_{\varphi} = (\neg_{\mathbb{A}^{\varphi}} x^{\varphi})_{\varphi}$
Join and meet	**Modal operators**
$(x^{\varphi})_{\varphi} \vee_{\mathfrak{A}} (y^{\varphi})_{\varphi} = (x^{\varphi} \vee_{\mathbb{A}^{\varphi}} y^{\varphi})_{\varphi}$	$\Diamond_{\mathfrak{A}}(x^{\varphi})_{\varphi} = (\Diamond_{\mathbb{A}^{\varphi}} x^{\varphi})_{\varphi}$
$(x^{\varphi})_{\varphi} \wedge_{\mathfrak{A}} (y^{\varphi})_{\varphi} = (x^{\varphi} \wedge_{\mathbb{A}^{\varphi}} y^{\varphi})_{\varphi}$	$\Box_{\mathfrak{A}}(x^{\varphi})_{\varphi} = (\Box_{\mathbb{A}^{\varphi}} x^{\varphi})_{\varphi}$

⊣

The product of a family of modal Boolean algebras $\{\mathbb{A}_i\}_{i \in I}$ is a modal Boolean algebra, where I is a (possibly uncountable) index set [34, Section 7]. Thus, $\mathfrak{A}$ is a modal Boolean algebra. Notice that one can also define the modal operators $\blacklozenge_{\mathfrak{A}}$ and $\blacksquare_{\mathfrak{A}}$ pointwise as follows: $\blacklozenge_{\mathfrak{A}}(x^{\varphi})_{\varphi} = \left(\blacklozenge_{\mathbb{A}^{\varphi}} x^{\varphi}\right)_{\varphi}$ and $\blacksquare_{\mathfrak{A}}(x^{\varphi})_{\varphi} = \left(\blacksquare_{\mathbb{A}^{\varphi}} x^{\varphi}\right)_{\varphi}$, for any $(x^{\varphi})_{\varphi} \in \mathfrak{A}$, such that $\blacklozenge_{\mathfrak{A}} \dashv \Box_{\mathfrak{A}}$ and $\Diamond_{\mathfrak{A}} \dashv \blacksquare_{\mathfrak{A}}$.

By Lemma 3.5.3 for every modal Boolean algebra $\mathbb{A}$ and for every $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ there is a refinement morphism g^{φ} from $\mathbb{A}$ to $\mathbb{A}^{\varphi}$. So it seems reasonable to expect that we can define a refinement morphism from $\mathbb{A}$ to $\mathfrak{A}$. And indeed we can.

Definition 3.5.6 *For every modal Boolean algebra $\mathbb{A}$, let the maps $G_{\mathbb{A}}$ and $F_{\mathbb{A}}$ be defined as follows:*

$$G_{\mathbb{A}} : \mathbb{A} \to \mathfrak{A}_{\mathbb{A}} \qquad\qquad F_{\mathbb{A}} : \mathfrak{A}_{\mathbb{A}} \to \mathbb{A}$$
$$x \mapsto (\prod_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} g^{\varphi})(x) \qquad\qquad ([h]^{\varphi})_{\varphi} \mapsto \bigvee_{\varphi} f^{\varphi}([h]^{\varphi})$$

where $g^{\varphi} : \mathbb{A} \to \mathbb{A}^{\varphi}$ and $f^{\varphi} : \mathbb{A}^{\varphi} \to \mathbb{A}$ are the maps defined in (3.11). ⊣

For the sake of readability and when it causes no confusion, we drop the subscripts and write G instead of $G_{\mathbb{A}}$ and F instead of $F_{\mathbb{A}}$.

The next lemma is an analogue of Lemma 3.5.3 and shows that $G_{\mathbb{A}}$ is indeed a refinement morphism and together with $F_{\mathbb{A}}$ forms an adjunction.

Lemma 3.5.7 *Let $\mathbb{A}$ be a modal Boolean algebra, and let $\mathfrak{A}_{\mathbb{A}}$ be the refinement algebra of $\mathbb{A}$. We have that*

1. *the map $G_{\mathbb{A}}$ is a refinement morphism,*
2. *the map $F_{\mathbb{A}}$ preserves $\mathbf{0}$ and finite joins,*
3. *$F_{\mathbb{A}} \dashv G_{\mathbb{A}}$.* ⊣

Proof *Item 1.* Let us show that $G_{\mathbb{A}}$ is a refinement morphism. Since for each $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, g^{φ} preserves $\mathbf{0}_{\mathbb{A}^{\varphi}}$ and $\vee_{\mathbb{A}^{\varphi}}$, it follows that $\prod_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} g^{\varphi} : \mathbb{A} \to \prod_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} \mathbb{A}^{\varphi}$ also satisfies those properties. It remains to prove that $\blacklozenge_{\mathfrak{A}} \circ G_{\mathbb{A}} \leq G_{\mathbb{A}} \circ \blacklozenge_{\mathbb{A}}$. It follows from Lemma 3.5.3 that each g^{φ} is a refinement morphism and satisfies (**AlgZag**); that is: $\blacklozenge_{\mathbb{A}^{\varphi}} \circ g^{\varphi} \leq g^{\varphi} \circ \blacklozenge_{\mathbb{A}}$, for every $\varphi \in \mathcal{L}_{\Box\alpha\forall}$. Let $x \in \mathbb{A}$, we have that

$$\begin{aligned}
\blacklozenge_{\mathfrak{A}} \circ G_{\mathbb{A}}(x) &= \blacklozenge_{\mathfrak{A}} \left(g^{\varphi}(x)\right)_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} && \text{(Def. } G\text{)}\\
&= \left(\blacklozenge_{\mathbb{A}^{\varphi}}(g^{\varphi})(x)\right)_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} && \text{(Def. } \blacklozenge_{\mathfrak{A}}\text{)}\\
&\leq \left(g^{\varphi}(\blacklozenge_{\mathbb{A}}(x))\right)_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} && (\blacklozenge_{\mathbb{A}^{\varphi}} \circ g^{\varphi} \leq g^{\varphi} \circ \blacklozenge_{\mathbb{A}})\\
&= (\prod_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} g^{\varphi})\,(\blacklozenge_{\mathbb{A}}(x))\\
&= G_{\mathbb{A}} \circ \blacklozenge_{\mathbb{A}}(x). && \text{(Def. } G\text{)}
\end{aligned}$$

Item 2. It is easy to see that $F_{\mathbb{A}}$ is monotone and preserves $\mathbf{0}$. We proceed to show that $F_{\mathbb{A}}$ preserves binary joins and then by induction we can easily prove that it preserves finite joins. Let $([h]^{\varphi})_{\varphi \in \mathcal{L}_{\Box\alpha\forall}}, ([k]^{\varphi})_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} \in \mathfrak{A}$. Then

$$\begin{aligned}
F_{\mathbb{A}}\Big(\big(([h]^{\varphi})_{\varphi} \vee ([k]^{\varphi})_{\varphi}\big)\Big) &= \bigvee_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} f^{\varphi}\Big(\big(([h]^{\varphi})_{\varphi} \vee ([k]^{\varphi})_{\varphi \in \mathcal{L}_{\Box\alpha\forall}}\big)\Big) && \text{(Def. } F_{\mathbb{A}}\text{)}\\
&= \bigvee_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} f^{\varphi}([h]^{\varphi})_{\varphi} \vee f^{\varphi}([k]^{\varphi})_{\varphi} && \text{(Lemma 3.5.3.2)}\\
&= \bigvee_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} f^{\varphi}([h]^{\varphi})_{\varphi} \vee \bigvee_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} f^{\varphi}([k]^{\varphi})_{\varphi} && (*)\\
&= F_{\mathbb{A}}\big(([h]^{\varphi})_{\varphi} \vee F_{\mathbb{A}}\big(([k]^{\varphi})_{\varphi}\big). && \text{(Def. } F_{\mathbb{A}}\text{)}
\end{aligned}$$

The equality marked $(*)$ holds because arbitrary joins distributes over finite joins.
Item 3 is immediate and follows from the fact that for each $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, $g^{\varphi} \dashv f^{\varphi}$. □

It is now time to introduce the algebraic semantics of refinement action model logic.

Definition 3.5.8 (Algebraic semantics of RAML) *Let $\mathcal{A} = \langle \mathbb{A}, V\rangle$ be an algebraic model. We define the* refinement *of $\mathcal{A}$ as the algebraic model $\mathcal{E} = (\mathfrak{A}, \mathcal{V})$ where $\mathfrak{A}$ is the refinement algebra of $\mathbb{A}$, and $\mathcal{V} : \mathsf{At} \to \mathfrak{A}$ where $\mathcal{V}(p) = (F_{\mathbb{A}} \circ V)(p)$. The extension map $[\![-]\!] : \mathcal{L}_{\Box\alpha\forall} \to \mathbb{A}$ of V over $\mathcal{L}_{\Box\alpha\forall}$ is defined recursively as in Def. 3.4.6, plus the following additional clause*

$$[\![\exists\varphi]\!]_{\mathcal{A}} := F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}).$$

We will need the next lemma to show the soundness of $\mathbb{RAML}$ with respect to the algebraic semantics.

Lemma 3.5.9 *For any formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, we have: $F_{\mathbb{A}}([\![\Diamond\varphi]\!]_{\mathcal{E}}) \leq \Diamond F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}})$.* ⊣

Proof Fix an algebraic model $\mathcal{A} = \langle \mathbb{A}, V\rangle$ and a formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$. We want to prove $F_{\mathbb{A}}([\![\Diamond\varphi]\!]_{\mathcal{E}}) \leq \Diamond F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}})$.

$$\begin{aligned}
F_{\mathbb{A}}([\![\Diamond\varphi]\!]_{\mathcal{E}}) &= \bigvee_{\gamma} f^{\gamma}([\![\Diamond\varphi]\!]_{\mathcal{A}^{\gamma}}) = \bigvee_{\gamma} \left([\![\langle\alpha^{\gamma}_{K^{\gamma}}\rangle\Diamond\varphi]\!]_{\mathcal{A}}\right) && \text{(Def. } F \text{ \& Lemma 3.5.3(4))}\\
&= \bigvee_{\gamma}\bigvee_{k\in K^{\gamma}} \left([\![Pre^{\gamma}_{\alpha}(k)]\!]_{\mathcal{A}} \wedge [\![\bigvee_{j\in R^{\gamma}_{\alpha}(k)} \Diamond\langle\alpha^{\gamma}_{j}\rangle\varphi]\!]_{\mathcal{A}}\right) && \text{(Lemma 3.4.7(6))}\\
&\leq \bigvee_{\gamma}\bigvee_{k\in K^{\gamma}} [\![\bigvee_{j\in R^{\gamma}_{\alpha}(k)} \Diamond\langle\alpha^{\gamma}_{j}\rangle\varphi]\!]_{\mathcal{A}} && (x\wedge y\leq x)\\
&= \bigvee_{\gamma}\bigvee_{k\in K^{\gamma}} \Diamond[\![\bigvee_{j\in R^{\gamma}_{\alpha}(k)} \langle\alpha^{\gamma}_{j}\rangle\varphi]\!]_{\mathcal{A}} && (\Diamond(\varphi\vee\psi) = \Diamond\varphi\vee\Diamond\psi)\\
&\leq \bigvee_{\gamma}\bigvee_{k\in K^{\gamma}} \Diamond[\![\bigvee_{k\in K^{\gamma}} \langle\alpha^{\gamma}_{k}\rangle\varphi]\!]_{\mathcal{A}} && (R^{\gamma}_{\alpha}(k)\subseteq K^{\gamma})\\
&= \bigvee_{\gamma}\bigvee_{k\in K^{\gamma}} \Diamond f^{\gamma}([\![\varphi]\!]_{\mathcal{A}^{\gamma}}) && \text{(Lemma 3.5.3(4))}\\
&= \bigvee_{\gamma} \Diamond f^{\gamma}([\![\varphi]\!]_{\mathcal{A}^{\gamma}}) = \Diamond\bigvee_{\gamma} f^{\gamma}([\![\varphi]\!]_{\mathcal{A}^{\gamma}})\\
&= \Diamond F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}) && \text{(Def. } F_{\mathbb{A}})
\end{aligned}$$

□

Theorem 3.5.10 (Soundness) *The axiomatisation $\mathbb{RAML}$ is sound with respect to the algebraic* **RAML**-*models.* ⊣

Proof We need to show that for any formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, if $\vdash_{\mathbb{RAML}} \varphi$ then $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$ for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V\rangle$. The proof is by well-founded induction on the structure of φ refined by refinement modal depth (maximum length of $\exists$ modalities binding each other) defined in the obvious way, i.e., we consider the order $<$ such that: if the refinement quantifier depth of φ is smaller than that of ψ, then $\varphi < \psi$, whereas if they are the same, then $\varphi < \psi$ if φ is a subformula of ψ. We must show that for all $\varphi \in \mathcal{L}_{B\Box\alpha}$, (For all $\psi < \varphi$: if $\vdash_{\mathbb{RAML}} \psi$ then for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V \rangle$, $[\![\psi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$)

$\Rightarrow$

(if $\vdash_{\mathbb{RAML}} \varphi$ then for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V \rangle$, $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$).

So, suppose that $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ and for every $\psi \in \mathcal{L}_{\Box\alpha\forall}$ with $\psi < \varphi$, if $\vdash_{\mathbb{RAML}} \psi$ then for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V \rangle$, $[\![\psi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$. We must show that if $\vdash_{\mathbb{RAML}} \varphi$ then for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V \rangle$, $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$. The soundness of axioms and rules of $\mathbb{AML}$ follows from the same reasoning used in [105] to show that $\mathbb{AML}$ is sound with respect to algebraic **AML**-models because the definition of $[\![-]\!]_{\mathcal{A}}$ (Def. 3.5.8) for the $\mathcal{L}_{\Box\alpha}$-fragment of $\mathcal{L}_{\Box\alpha\forall}$ is identical to the algebraic semantics defined in [105]. So we only need to consider the axioms and rules of $\mathbb{RML}$.

- Rule **NecR**. Suppose that $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ and for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V \rangle$, $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$. We will show that $[\![\forall\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$, for all algebraic models $\mathcal{A} = \langle \mathbb{A}, V \rangle$. Let $\mathcal{A} = \langle \mathbb{A}, V \rangle$ be an algebraic model. Since $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$, we have that $[\![\neg\varphi]\!]_{\mathcal{A}} = \mathbf{0}_{\mathbb{A}}$. Then we have

$$[\![\forall\varphi]\!]_{\mathcal{A}} = [\![\neg\exists\neg\varphi]\!]_{\mathcal{A}} = \neg F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{A}}) = \neg F_{\mathbb{A}}([\![\mathbf{0}_{\mathbb{A}}]\!]_{\mathcal{A}})$$

By Lemma 3.5.7(2), we have that $F_{\mathbb{A}}(\mathbf{0}_{\mathbb{A}}) = \mathbf{0}_{\mathbb{A}}$. Hence, $[\![\forall\varphi]\!]_{\mathcal{A}} = \neg\mathbf{0}_{\mathbb{A}} = \mathbf{1}_{\mathbb{A}}$.

- Axiom **RProp**. It suffices to show $[\![\exists p]\!]_{\mathcal{A}} = [\![p]\!]_{\mathcal{A}}$ and $[\![\exists\neg p]\!]_{\mathcal{A}} = [\![\neg p]\!]_{\mathcal{A}}$. Let $p \in \mathsf{At}$. First of all, by Def. 3.4.4 we have

$$[\![p]\!]_{\mathcal{A}^\varphi} = [\prod_{a^\varphi} V(p)].$$

Then for all formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ and by the definition of f^φ, we have

$$f^\varphi([\![p]\!]_{\mathcal{A}^\varphi}) = \bigvee_{k \in K^\varphi} \prod_{a^\varphi} V(p)(k) \wedge Pre^\varphi_a(k) = \bigvee_{k \in K^\varphi} V(p) \wedge Pre^\varphi_a(k).$$

Hence, $f^\varphi([\![p]\!]_{\mathcal{A}^\varphi}) \leq V(p)$. But since φ was arbitrary, it follows by the definition of least upper bound that

$$\bigvee_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} f^\varphi([\![p]\!]_{\mathcal{A}^\varphi}) \leq [\![p]\!]_{\mathcal{A}}.$$

So, by the definition of $\mathcal{V}$, and the definition of $F_{\mathbb{A}}$, we obtain that

$$F_{\mathbb{A}}([\![p]\!]_{\mathcal{E}}) \leq [\![p]\!]_{\mathcal{A}}.$$

For the other direction, according to the construction of multi-pointed action model $\alpha^p_{K^p}$ for the atomic proposition p [84, Lemma V.2], $[\![\langle\alpha^p_{K^p}\rangle p]\!]_{\mathcal{A}} = [\![p]\!]_{\mathcal{A}}$. This implies that $[\![p]\!]_{\mathcal{A}} = f^p([\![p]\!]_{\mathcal{A}^p})$ and $[\![p]\!]_{\mathcal{A}} \leq \bigvee_{\varphi \in \mathcal{L}_{\Box\alpha\forall}} f^\varphi([\![p]\!]_{\mathcal{A}^\varphi}) = F_{\mathbb{A}}([\![p]\!]_{\mathcal{E}})$, as required. The other identity for $\neg p$ can be proved in a similar way.

- Axiom **R**. We need to show for a given algebraic model $\mathcal{A} = \langle \mathbb{A}, V \rangle$, $[\![\forall(\varphi \to \psi)]\!]_{\mathcal{A}} \leq [\![\forall\varphi \to \forall\psi]\!]_{\mathcal{A}}$. Since $\forall := \neg\exists\neg$, we observe that

$$[\![\forall(\varphi \to \psi)]\!]_{\mathcal{A}} = [\![\neg\exists\neg(\varphi \to \psi)]\!]_{\mathcal{A}} = \neg F_{\mathbb{A}}([\![\neg(\varphi \to \psi)]\!]_{\mathcal{E}})$$

and

$$[\![\forall\varphi \to \forall\psi]\!]_{\mathcal{A}} = [\![\exists\neg\varphi \vee \neg\exists\neg\psi]\!]_{\mathcal{A}} = F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \vee \neg F_{\mathbb{A}}([\![\neg\psi]\!]_{\mathcal{E}}).$$

Hence, it is enough to show that

$$\neg F_{\mathbb{A}}([\![\neg(\varphi \to \psi)]\!]_{\mathcal{E}}) \leq F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \vee \neg F_{\mathbb{A}}([\![\neg\psi]\!]_{\mathcal{E}})$$

Note that since $\neg\varphi \vee \neg\psi \leftrightarrow \neg\varphi \vee \neg(\varphi \to \psi)$, it follows that $[\![\neg\varphi \vee \neg\psi]\!]_{\mathcal{E}} = [\![\neg\varphi \vee \neg(\varphi \to \psi)]\!]_{\mathcal{E}}$. Hence, $F_{\mathbb{A}}([\![\neg\varphi \vee \neg\psi]\!]_{\mathcal{E}}) = F_{\mathbb{A}}([\![\neg\varphi \vee \neg(\varphi \to \psi)]\!]_{\mathcal{E}}))$. But $F_{\mathbb{A}}$ preserves $\vee$, hence

$$F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \vee F_{\mathbb{A}}([\![\neg\psi]\!]_{\mathcal{E}}) = F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \vee F_{\mathbb{A}}([\![\neg(\varphi \to \psi)]\!]_{\mathcal{E}})).$$

By applying negation on both sides we get

$$\neg F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \wedge \neg F_{\mathbb{A}}([\![\neg\psi]\!]_{\mathcal{E}}) = \neg F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \wedge \neg F_{\mathbb{A}}([\![\neg(\varphi \to \psi)]\!]_{\mathcal{E}})$$

And this implies that

$$\neg F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \wedge \neg F_{\mathbb{A}}([\![\neg(\varphi \to \psi)]\!]_{\mathcal{E}}) \leq \neg F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \wedge \neg F_{\mathbb{A}}([\![\neg\psi]\!]_{\mathcal{E}}) \tag{3.14}$$

Using that in every Boolean algebra, $x \wedge y \leq x \wedge z$ implies that $y \leq \neg x \vee z$, (3.14) implies that $\neg F_{\mathbb{A}}([\![\neg(\varphi \to \psi)]\!]_{\mathcal{E}}) \leq F_{\mathbb{A}}([\![\neg\varphi]\!]_{\mathcal{E}}) \vee \neg F_{\mathbb{A}}([\![\neg\psi]\!]_{\mathcal{E}})$.

- Axiom **RK**. We need to show that for every algebraic model $\mathcal{A} = \langle \mathbb{A}, V\rangle$, and all $\Phi \subseteq \mathcal{L}_{\Box\alpha\forall}$, $[\![\exists\nabla\Phi]\!]_{\mathcal{A}} = [\![\bigwedge \Diamond\exists\Phi]\!]_{\mathcal{A}}$. We first show $[\![\exists\nabla\Phi]\!]_{\mathcal{A}} \leq [\![\bigwedge \Diamond\exists\Phi]\!]_{\mathcal{A}}$.

$$\begin{aligned}
[\![\exists\nabla\Phi]\!]_{\mathcal{A}} &= F_{\mathbb{A}}\left([\![\Box\left(\bigvee_{\varphi\in\Phi}\varphi\right) \wedge \bigwedge_{\varphi\in\Phi}\Diamond\varphi]\!]_{\mathcal{E}}\right) \\
&\leq F_{\mathbb{A}}\left([\![\bigwedge_{\varphi\in\Phi}\Diamond\varphi]\!]_{\mathcal{E}}\right) \leq \bigwedge_{\varphi\in\Phi} F_{\mathbb{A}}([\![\Diamond\varphi]\!]_{\mathcal{E}}) && \text{(monotony of } F_{\mathbb{A}}\text{)} \\
&\leq \bigwedge_{\varphi\in\Phi}\Diamond F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}) = [\![\bigwedge\Diamond\exists\Phi]\!]_{\mathcal{A}}. && \text{(Lemma 3.5.9)}
\end{aligned}$$

To show the other inequality, i.e., $[\![\bigwedge \Diamond\exists\Phi]\!]_{\mathcal{A}} \leq [\![\exists\nabla\Phi]\!]_{\mathcal{A}}$, we first prove that

$$\bigwedge_{\varphi\in\Phi}\Diamond F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}) \leq f^{\nabla\Phi}([\![\nabla\Phi]\!]_{\mathcal{A}^{\nabla\Phi}}). \tag{3.15}$$

Let $\alpha^{\nabla\Phi}$ be the action model constructed by Hales as shown in Theorem 3.2.19. By the definition of $f^{\nabla\Phi}$ and the structure of $\alpha^{\nabla\Phi}$, we have

$$\begin{aligned}
f^{\nabla\Phi}([\![\nabla\Phi]\!]_{\mathcal{A}^{\nabla\Phi}}) &= \bigvee_{k\in K^{\nabla\Phi}} [\![\langle\alpha_k^{\nabla\Phi}\rangle\nabla\Phi]\!]_{\mathcal{A}} && \text{(Lemma 3.5.3(4))} \\
&= [\![\langle\alpha_{k^\star}^{\nabla\Phi}\rangle\nabla\Phi]\!]_{\mathcal{A}} \vee \bigvee_{\varphi\in\Phi}\bigvee_{k\in K^{\varphi}} [\![\langle\alpha_k^{\nabla\Phi}\rangle\nabla\Phi]\!]_{\mathcal{A}}.
\end{aligned}$$

Therefore,

$$[\![\langle\alpha_{k^\star}^{\nabla\Phi}\rangle\nabla\Phi]\!]_{\mathcal{A}} \leq f^{\nabla\Phi}([\![\nabla\Phi]\!]_{\mathcal{A}^{\nabla\Phi}}). \tag{3.16}$$

Also, by the definition of $\nabla\Phi$ and Lemma 3.4.7, we obtain that

$$[\![\langle\alpha_{k^\star}^{\nabla\Phi}\rangle\nabla\Phi]\!]_{\mathcal{A}} = [\![\langle\alpha_{k^\star}^{\nabla\Phi}\rangle(\Box\bigvee_{\varphi\in\Phi}\varphi \wedge \bigwedge_{\varphi\in\Phi}\Diamond\varphi)]\!]_{\mathcal{A}}$$

$$= [\![\langle \alpha_{k^*}^{\nabla\Phi} \rangle \Box \bigvee_{\varphi\in\Phi} \varphi]\!]_{\mathcal{A}} \wedge [\![\langle \alpha_{k^*}^{\nabla\Phi} \rangle \bigwedge_{\varphi\in\Phi} \Diamond\varphi]\!]_{\mathcal{A}}.$$

Moreover, for every pointed action model $\alpha_k = (\langle K, R_\alpha, Pre_\alpha\rangle, k)$ over $\mathbb{A}$, a simple semantic argument shows that

$$[\![\langle \alpha_k\rangle\gamma]\!]_{\mathcal{A}} = [\![Pre_\alpha(k)]\!]_{\mathcal{A}} \wedge [\![[\alpha_k]\gamma]\!]_{\mathcal{A}} \tag{3.17}$$

So we have that

$$[\![\langle \alpha_{k^*}^{\nabla\Phi} \rangle \Box \bigvee_{\varphi\in\Phi} \varphi]\!]_{\mathcal{A}} = [\![Pre_\alpha^{\nabla\Phi}(k^*)]\!]_{\mathcal{A}} \wedge [\![[\alpha_{k^*}^{\nabla\Phi}] \Box \bigvee_{\varphi\in\Phi} \varphi]\!]_{\mathcal{A}}$$

and

$$[\![\langle \alpha_{k^*}^{\nabla\Phi} \rangle \bigwedge_{\varphi\in\Phi} \Diamond\varphi]\!]_{\mathcal{A}} = [\![Pre_\alpha^{\nabla\Phi}(k^*)]\!]_{\mathcal{A}} \wedge [\![[\alpha_{k^*}^{\nabla\Phi}] \bigwedge_{\varphi\in\Phi} \Diamond\varphi]\!]_{\mathcal{A}}.$$

By [84, Lemma V.2], it follows that $\vdash_{\mathbb{RAML}} [\alpha_{k^*}^{\nabla\Phi}]\Box \bigvee_{\varphi\in\Phi}\varphi$ and $\vdash_{\mathbb{RAML}} [\alpha_{k^*}^{\nabla\Phi}] \bigwedge_{\varphi\in\Phi}\Diamond\varphi$. Then since the refinement modal depth of $[\alpha_{k^*}^{\nabla\Phi}]\Box \bigvee_{\varphi\in\Phi}\varphi$ and $[\alpha_{k^*}^{\nabla\Phi}] \bigwedge_{\varphi\in\Phi}\Diamond\varphi$ is one less than that of $\bigwedge_{\varphi\in\Phi}\Diamond\exists\varphi$, we can now use induction and it follows that $[\![[\alpha_{k^*}^{\nabla\Phi}]\Box \bigvee_{\varphi\in\Phi}\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$ and $[\![[\alpha_{k^*}^{\nabla\Phi}] \bigwedge_{\varphi\in\Phi}\Diamond\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$. Hence, we obtain that

$$[\![\langle \alpha_{k^*}^{\nabla\Phi} \rangle \nabla\Phi]\!]_{\mathcal{A}} = [\![Pre_\alpha^{\nabla\Phi}(k^*)]\!]_{\mathcal{A}} = [\![\bigwedge_{\varphi\in\Phi} \Diamond\exists\varphi]\!]_{\mathcal{A}} = \bigvee_{\varphi\in\Phi} F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}). \tag{3.18}$$

The above identity together with (3.16) implies (3.15) holds, i.e.,

$$\bigwedge_{\varphi\in\Phi} F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}) \leq f^{\nabla\Phi}([\![\nabla\Phi]\!]_{\mathcal{A}^{\nabla\Phi}}).$$

It then follows from the definition of $F_{\mathbb{A}}$ that

$$\begin{aligned} \bigwedge_{\varphi\in\Phi} \Diamond F_{\mathbb{A}}([\![\varphi]\!]_{\mathcal{E}}) &\leq f^{\nabla\Phi}([\![\nabla\Phi]\!]_{\mathcal{A}^{\nabla\Phi}}) \\ &\leq \bigvee_{\varphi\in\mathcal{L}_{\Box\alpha\forall}} f^{\varphi}([\![\nabla\Phi]\!]_{\mathcal{A}^{\nabla\Phi}}) \\ &= F_{\mathbb{A}}([\![\nabla\Phi]\!]_{\mathcal{E}}). \end{aligned} \tag{3.15}$$

Therefore, $[\![\exists\nabla\Phi]\!]_{\mathcal{A}} = [\![\bigwedge_{\varphi\in\Phi}\Diamond\exists\varphi]\!]_{\mathcal{A}}$, as desired.

□

Theorem 3.5.11 (Completeness) *The axiomatisation $\mathbb{RAML}$ is complete with respect to the algebraic **RAML**-models, i.e., for every formula $\varphi \in \mathcal{L}_{\Box\alpha\forall}$, and for every algebraic model $\mathcal{A} = \langle \mathbb{A}, V\rangle$, if $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$ then $\vdash_{\mathbb{RAML}} \varphi$.* ⊣

Proof Let $\varphi \in \mathcal{L}_{\Box\alpha\forall}$ be such that $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$ for every algebraic model $\mathcal{A} = \langle \mathbb{A}, V\rangle$. Then there is a formula ψ in the sublanguage $\mathcal{L}_{\Box}$ that does not contain any refinement quantifier [31, Prop. 36] or action model quantifiers [49] such that $\vdash_{\mathbb{RAML}} \varphi \leftrightarrow \psi$. By soundness, we have that $[\![\varphi \leftrightarrow \psi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$ and since $[\![\varphi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$, it follows that $[\![\psi]\!]_{\mathcal{A}} = \mathbf{1}_{\mathbb{A}}$. From the completeness of **K** with respect to algebraic models [25], we obtain that ψ is a theorem in **K**. This implies that φ is a theorem in **RAML** [31, Prop. 37]. □

3.6 Conclusion and future work

We have proposed an algebraic semantics for refinement action modal logic. Using action model synthesis and the algebraic characterisation of epistemic updates, we have introduced the abstract notion of refinement on modal Boolean algebras, and showed the soundness and completeness of $\mathbb{RAML}$ with respect to this algebraic semantics.

Our methodology builds on and further develops recent work [105, 114] applying duality theory to dynamic epistemic logic. As part of this research program, proof systems for intuitionistic **AML** have been introduced [80, 71], and gave rise to the novel methodology of multi-type display calculi [70], which has been applied not only to **AML** [72], but also to propositional dynamic logic [69] and inquisitive logic [73]. A natural direction is to pursue this research program also on refinement modal logic. We plan to weaken the classical propositional modal logical base to a non-classical propositional modal logical base, and to develop multi-type calculi for such non-classical modal logics with refinement quantifiers, for example refinement intuitionistic (modal) logic.

Another step to take would be to generalise the algebraic semantics of **RAML** to the multi-agent framework. In this framework, the refinement modality $\exists$ is indexed by an agent, hence we have modalities $\{\exists_i\}_{i \in Ag}$ where Ag is the set of agents. The only difficulty in generalising our result is to prove, algebraically, the soundness of the additional axioms:

$$\exists_i \nabla_j \Phi \leftrightarrow \nabla_j \{\exists_i \varphi\}_{\varphi \in \Phi} \qquad \text{where } i \neq j$$
$$\exists_i \bigwedge_{j \in J} \nabla_j \Phi^j \leftrightarrow \bigwedge_{j \in J} \exists_i \nabla_j \Phi^j \qquad \text{where } J \subseteq Ag.$$

Indeed, the soundness proofs can be quite involved, as the reader can see in Section 3.5. A good proof system for the multi-agent refinement modal logic would be useful.

4

Bilattice dynamic epistemic logic

Contents

You are lost on Place Stanislas in the historical center of Nancy and you need to catch a train. So you accost a friendly and French looking person and there you go, pointing to the right: "Is this the way to the railway station?" "Oui." (Yes.) Merci, etc., you each go your way, but, a few moments later, while still remaining in some doubt, you ask another person, and then pointing in the opposite direction: "Is this the way to the railway station?" "Oui." What will you do? (First lesson: when asking directions, never *suggestively point in one direction.) You will probably resolve the inconsistency by yet further communication (or consultation of a map, say) before you continue on your way. And sure enough, the next person you ask does not even answer the question and shrugs her shoulders before walking on. Inconsistent or absent responses in dynamic interaction are just as common as inconsistency in static information. Propositions that can be true, false, both (true and false), or neither are modelled with bilattices. In this work we investigate the dynamic modal logic of bilattices, where not only propositions but also actions have four-valued features.*

4.1 Introduction

In the past decades, reasoning about *knowledge* and *information change* has gained a prominent place in various areas of artificial intelligence and computer science such as distributed systems [87], protocol verification [88], and game theory [7]. In these areas agents have to deal with *incomplete* and *inconsistent* information, and by incomplete information we mean lacking or missing information. For example, in distributed systems, agents receive information from multiple sources that may be inconsistent. Moreover, in real-world situations, agents do not have complete information about all aspects of the world and their reasoning power is bounded by thresholds such as time and limited memory [55]. Under such circumstances, applying a classical approach to model information change may not be appropriate because it suffers from the *logical omniscience* problem [154]; that is, the agents know all the consequences of what they know. As a result, they cannot hold contradictory knowledge without "knowing" every sentence of the language, because a contradiction classically entails any formula. Several approaches have been proposed to formalise inconsistent and incomplete information in the literature, see e.g. [15, 16, 108, 106, 55, 57]. To set the stage for future discussion, we are going to review the most closely related works. In [15] Belnap proposed a *four-valued logic* whose semantics involves, besides the classical truth values t and f, two intermediate values: $\top$ (both true and false) for handling inconsistent information and $\bot$ (neither true nor false) for incomplete information. In this logic, each atomic formula can be assigned one of the four values chosen from the set $\mathbf{4} = \{\mathsf{t}, \mathsf{f}, \bot, \top\}$. Belnap observed that his four values can be arranged in a lattice in two ways: ordering them either by information degree (the knowledge order $\leq_k$) or by the truth degree (the truth order $\leq_t$). The set $\mathbf{4}$ together with $\leq_k$ and $\leq_t$ forms two complete, distributive lattices, which are shown in Figure 4.1.

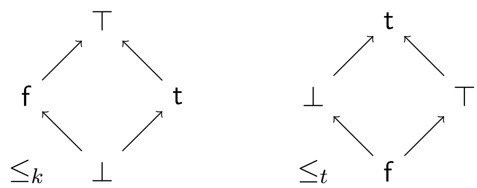

Figure 4.1. The four-element Belnap lattice in its two orders, the bilattice FOUR

Given two truth values x and y, $x \leq_t y$ can be read as "y is at least as true as x", while $x \leq_k y$ means that "y contains at least as much information as x".

Belnap's four-valued logic inspired Levesque to address the logical omniscience problem. In [108] he proposed a *logic of explicit and implicit belief*. Explicit beliefs are actively entertained by the agent, whereas implicit beliefs include the logical consequences of her explicit beliefs. This logic has a modality for explicit belief and a modality for implicit belief. The interpretation of these modalities is based on *situation semantics* [14]. Unlike in possible worlds, in a situation a sentence can be true, false, both true and false (*incoherent situation*), or neither true nor false (*incomplete situation*). From our perspective, [108] establishes a significant link between many-valued logics and epistemic logics. An objection raised against Levesque's model is that it is restricted to a single agent environment and therefore does not account for nested beliefs [136]. Fagin and Halpern address multi-agent belief in their logic of knowledge (or belief) and awareness [57]. The semantics of this awareness logic is based on possible worlds and does not allow the agents to have contradictory knowledge, but the awareness function at each possible world provides an effect that is similar to an incomplete situation. In [149, 148], Sim compares the approaches of [108] and [57] in detail and shows that the situations of [108] and the Kripke models with (un)awareness of [57] can be associated with a model based on a *bilattice* structure.

Bilattices are algebraic structures introduced by Ginsberg [77] to unify logical formalisms for default reasoning and non-monotonic reasoning. A bilattice is a set B equipped with two partial orders, the knowledge order ($\leq_k$) and the truth order ($\leq_t$), such that $(B, \leq_k)$ and $(B, \leq_t)$ are both complete lattices. The partial orders $\leq_k$ and $\leq_t$ have similar interpretations as in Belnap's logic. Belnap's four-element lattice is the smallest non-trivial bilattice. It is called FOUR (see Figure 4.1).

Bilattices have found applications in different research areas such as logic programming [66], semantics of natural language questions [126] and philosophical logic [65, 68]. In the 1990s, Arieli and Avron [5, 6] carried bilattices to a new stage introducing bilattice-based logical systems that are suitable for non-monotonic and paraconsistent reasoning. Later on, Jung and Rivieccio [98] introduced a modal expansion of the logic of [5] that can be used to reason about knowledge, belief, time, and obligation. The formulas of this logic are interpreted in *four-valued Kripke models* wherein both the accessibility relations and the valuations are four-valued. Four-valued accessibility relations go back to Fitting [66, 67], who suggested a family of many-valued modal logics and generalised Kripke models involving many-valued accessibility relations. He argued in [67] that many-valued accessibility relations are natural to formalise that some worlds alternative to the real world are more relevant than others.

A similar formalism to that of [98] was proposed by Odintsov and Wansing [127]. They studied a Belnapian version of the basic modal logic **K**. The semantics of this logic is based on Kripke models where valuations are four-valued (as in [98]), however, the accessibility relation is two-valued. Because of this, the modal operators of [127] differ from those of [98], although the propositional base of both logics is the same. The formalism of [98] is more general, because one can define the modal operators of [127] in the language of [98], but not the other way round [98, Prop. 2].

In this chapter we develop a bilattice-based modal logic with dynamic operators that enable us to reason about information change in the presence of incomplete and inconsistent information. We build the *bilattice action model logic* (**BAML**) by combining the *action model logic* (**AML**) of [11] with the bilattice-valued modal logic of [98]. The logic **AML** extends basic modal logic with an operator for reasoning about the effects of epistemic actions, as represented by action models. Epistemic actions are events by which agents receive new information about the world, whilst leaving the facts of the world itself unchanged. An action model is a relational structure similar to a Kripke model, where the accessibility relation between two actions (points in the action model domain) represents an agent's uncertainty as to which action actually occurred. The structure of action models should of course fit that of Kripke models, with four-valued accessibility relations. How to give intuitive interpretations to such four-valued action models is non-trivial, and we will give this ample attention. Formally, epistemic changes are modeled via the so-called *product update* construction on the Kripke models that provides a relational semantics for **AML**. Through the product update, a Kripke model encoding the current epistemic setup of a group of agents is replaced by an updated model.

An adequate formal treatment of **AML** and dynamic epistemic logics, from a syntactic as well as a semantic point of view, faces non-trivial technical problems which become even more serious when moving to a non-classical setting [71]. Such problems can be addressed in an algebraic framework. An elegant and versatile approach to the algebraic treatment of dynamic epistemic logic has been developed in a recent series of papers [105, 114, 138, 139, 35], in which the authors define non-classical counterparts of dynamic logics. In particular, the method from [114] has been applied by Rivieccio [138, 139] to obtain an algebraic semantics of bilattice-based public announcement logic, together with a sound and complete axiomatisation.

The contributions of this chapter are as follows: We extend the results of [138, 139] to bilattice action model logic (**BAML**). On the algebraic side, this extension consists of generalising the product update construction on modal bilattices from public announcements to arbitrary epistemic updates given by action models. The technical development relies on the methods from [105] and the correspondence between modal bilattices and bimodal Boolean algebras via the twist structure representation [137]. On the model side, this correspondence dualises to one between Kripke models with a four-valued relation and bimodal Kripke models (i.e. Kripke models with two two-valued relations). As a second contribution, we provide a Hilbert system for **BAML**, which is sound and complete with respect to the algebraic semantics. We restrict ourselves to the single-agent setting, but the multi-agent generalisation of our framework is straightforward. A final contribution consists of some motivating examples for bilattice-based dynamic epistemic logics such as **BAML**. Such examples have so far been missing in the literature.

The chapter is organised as follows. Section 4.2 recalls the necessary definitions and results on bilattice modal logic. It describes the static modal fragment on which we build our bilattice-based dynamic epistemic logic. Section 4.3 expounds the technical details of the update mechanism on the algebraic structures (modal bilattices), and introduces an algebraic semantics and a relational semantics for our logic. In Section 4.4 we introduce a Hilbert-style calculus for **BAML**, and we show its soundness and completeness. Completeness is shown by a reduction to the static fragment. Section 4.5 gives a detailed case study illustrating the usage of epistemic dynamics in a bilattice setting. Readers wishing to sharpen their intuitions on knowledge (change) and bilattices, or wanting to ascertain the relevance of our framework for such settings, are suggested to read this section earlier.

4.2 Bilattice modal logic

In this section we recall basic definitions and facts about bilattice modal logic, mainly from [5, 98], that will be needed to develop our bilattice-based action model logic. We refer the reader to [98, 140] for further details, as well as for background discussion and motivation on bilattices (see also Section 4.5).

4.2.1 Propositional logic of bilattices

The non-modal, propositional base of bilattice modal logic is the four-valued logic introduced by Arieli and Avron [5], which can be defined using Belnap's four-element lattice FOUR (Figure 4.1). In this logic, FOUR is viewed as an algebra having operations $\langle \wedge, \vee, \otimes, \oplus, \supset, \sim, \mathsf{f}, \mathsf{t}, \bot, \top \rangle$ of type $\langle 2, 2, 2, 2, 2, 1, 0, 0, 0, 0 \rangle$ such that both reducts $\langle \mathsf{FOUR}, \wedge, \vee, \mathsf{f}, \mathsf{t} \rangle$ and $\langle \mathsf{FOUR}, \otimes, \oplus, \bot, \top \rangle$ are bounded distributive lattices, where the lattice orders are denoted, respectively, by $\leq_t$ (*truth order*) and $\leq_k$ (*knowledge order*). The truth table for the operations $\wedge$ and $\vee$ is given as follows:

$\wedge$	t	$\top$	$\bot$	f
t	t	$\top$	$\bot$	f
$\top$	$\top$	$\top$	f	f
$\bot$	$\bot$	f	$\bot$	f
f	f	f	f	f

$\vee$	t	$\top$	$\bot$	f
t	t	t	t	t
$\top$	t	$\top$	t	$\top$
$\bot$	t	t	$\bot$	$\bot$
f	t	$\top$	$\bot$	f

The *binary weak implication operation* $\supset$ is defined by

$$x \supset y = \begin{cases} y & \text{if } x \in \{\mathsf{t}, \top\}, \\ \mathsf{t} & \text{if } x \notin \{\mathsf{t}, \top\}. \end{cases}$$

The truth table of $\supset$ is given in Table 4.1.

The *bilattice negation* is a unary operation $\sim$ having $\bot$ and $\top$ as fixed points and such that $\sim\mathsf{f} = \mathsf{t}$ and $\sim\mathsf{t} = \mathsf{f}$. The truth table of the bilattice negation $\sim$ is given in Table 4.1. The operations $\otimes$ and $\oplus$ need not be included in the primitive signature because they can be defined as terms in the language over $\langle \wedge, \vee, \supset, \sim, \mathsf{f}, \mathsf{t}, \bot, \top \rangle$. We define them here together with a number of other operations:

$$\begin{array}{llll} x \otimes y := (x \wedge \bot) \vee (y \wedge \bot) \vee (x \wedge y) & & x \oplus y := (x \wedge \top) \vee (y \wedge \top) \vee (x \wedge y) \\ x \to y := (x \supset y) \wedge (\sim y \supset \sim x) & & x \Leftrightarrow y := (x \supset y) \wedge (y \supset x) \\ x \leftrightarrow y := (x \to y) \wedge (y \to x). & & \neg x := x \supset \mathsf{f} \\ x * y := \neg(y \to \neg x). & & \end{array}$$

Table 4.1. Bilattice operations.

The operation $\neg$ provides an alternative negation (that one might call *two-valued (classical) negation*, to distinguish it from the bilattice negation $\sim$; note that $\neg x$ only takes values t and f). We note that our notation for bilattice negation and two-valued negation is different from that of [138, 139] where the bilattice negation is denoted by $\neg$ and two-valued negation is denoted by $\sim$. The operation $\to$ is an alternative implication called *strong implication*, which is adjoint to the operation $*$ with respect to the truth order $\leq_t$, called *strong conjunction* or *fusion*. The truth tables of the bilattice operations in FOUR are displayed below:

$\otimes$	t	$\top$	$\bot$	f
t	t	t	$\bot$	$\bot$
$\top$	t	$\top$	$\bot$	f
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
f	$\bot$	f	$\bot$	f

$\oplus$	t	$\top$	$\bot$	f
t	t	$\top$	t	$\top$
$\top$	$\top$	$\top$	$\top$	$\top$
$\bot$	t	$\top$	$\bot$	f
f	$\top$	$\top$	$\bot$	f

$\to$	t	$\top$	$\bot$	f
t	t	f	$\bot$	f
$\top$	t	$\top$	$\bot$	f
$\bot$	t	$\bot$	t	$\bot$
f	t	t	t	t

$\leftrightarrow$	t	$\top$	$\bot$	f
t	t	f	$\bot$	f
$\top$	f	$\top$	$\bot$	f
$\bot$	$\bot$	$\bot$	t	$\bot$
f	f	f	$\bot$	t

	t	$\top$	$\bot$	f
$\sim$	f	$\top$	$\bot$	t

	t	$\top$	$\bot$	f
$\neg$	f	f	t	t

$*$	t	$\top$	$\bot$	f
t	t	t	$\bot$	f
$\top$	t	$\top$	$\bot$	f
$\bot$	$\bot$	$\bot$	f	f
f	f	f	f	f

$\Leftrightarrow$	t	$\top$	$\bot$	f
t	t	$\bot$	t	t
$\top$	$\top$	$\top$	$\bot$	f
$\bot$	$\bot$	$\bot$	t	t
f	f	f	t	t

$\supset$	t	$\top$	$\bot$	f
t	t	$\top$	$\bot$	f
$\top$	t	$\top$	$\bot$	f
$\bot$	t	t	t	t
f	t	t	t	t

Table 4.2. The truth table of bilattice operations in FOUR.

The logic of bilattices of Arieli and Avron [5] can then be introduced as the propositional logic defined by the pair $\langle \mathsf{FOUR}, \{\mathsf{t}, \top\} \rangle$ as follows. Let At be a countable set of atomic propositions. The language $\mathcal{L}_B$ is defined by the following grammar:

$$\mathcal{L}_B \ni \varphi ::= \mathsf{f} \mid \mathsf{t} \mid \top \mid \bot \mid p \in \mathsf{At} \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \supset \varphi \mid \sim\varphi.$$

A valuation $V : \mathsf{At} \to \mathsf{FOUR}$ is a function that assigns a truth value from FOUR to each atomic proposition. Every valuation V has the unique extension $[\![.]\!]\colon \mathcal{L}_B \to \mathsf{FOUR}$ which is defined in a standard way. Given subsets $\Gamma, \{\varphi\} \subseteq \mathcal{L}_B$, we say φ is entailed by Γ (notation: $\Gamma \models_{\mathbf{LB}} \varphi$), if for all valuations $V\colon \mathsf{At} \to \mathsf{FOUR}$, $[\![\gamma]\!] \in \{\mathsf{t}, \top\}$ for all $\gamma \in \Gamma$ implies $[\![\varphi]\!] \in \{\mathsf{t}, \top\}$. We call a formula $\varphi \in \mathcal{L}_B$ a *tautology* , if $\models_{\mathbf{LB}} \varphi$. *The logic of bilattices* (**LB**) is defined as the set of formulas $\varphi \in \mathcal{L}_B$ such that φ is a tautology. Arieli and Avron [5] provided an axiomatisation $\mathbb{LB}$ for **LB**, which is given in Table 4.3. The axioms of $\mathbb{LB}$ are the axioms of propositional logic in the language $\langle \wedge, \vee, \supset, \mathsf{f}, \mathsf{t}\rangle$, plus the axioms that characterise the interaction of negation with other operations and constants.

($\supset$**1**)	$\varphi \supset (\psi \supset \varphi)$	($\supset$**f**)	$\mathsf{f} \supset \varphi$
($\supset$**2**)	$(\varphi \supset (\psi \supset \chi)) \supset ((\varphi \supset \psi) \supset (\varphi \supset \chi))$	($\supset\bot$)	$\bot \supset \varphi$
($\supset$**3**)	$((\varphi \supset \psi) \supset \varphi) \supset \varphi$	($\supset\top$)	$\varphi \supset \top$
($\sim\sim$)	$\varphi \Leftrightarrow {\sim}{\sim}\varphi$	($\supset$**t**)	$\varphi \supset \mathsf{t}$
($\wedge\supset$)	$(\varphi \wedge \psi) \supset \varphi \qquad (\varphi \wedge \psi) \supset \psi$	($\sim\vee$)	${\sim}(\varphi \vee \psi) \Leftrightarrow ({\sim}\varphi \wedge {\sim}\psi)$
($\supset\wedge$)	$\varphi \supset (\psi \supset (\varphi \wedge \psi))$	($\sim\supset$)	${\sim}(\varphi \supset \psi) \Leftrightarrow (\varphi \wedge {\sim}\psi)$
($\sim\wedge$)	${\sim}(\varphi \wedge \psi) \Leftrightarrow ({\sim}\varphi \vee {\sim}\psi)$		
($\supset\vee$)	$\varphi \supset (\varphi \vee \psi) \qquad \psi \supset (\varphi \vee \psi)$	(**MP**)	from φ and $\varphi \supset \psi$ infer ψ
($\vee\supset$)	$(\varphi \supset \chi) \supset ((\psi \supset \chi) \supset ((\varphi \vee \psi) \supset \chi))$		

Table 4.3. The proof system $\mathbb{LB}$.

4.2.2 Bilattice modal logic: relational semantics

In this subsection, we recall the syntax and relational semantics of bilattice modal logic of [98].

Definition 4.2.1 (Syntax of $\mathcal{L}_{B\Box}$) *Let* At *be a countable set of atomic propositions. The language $\mathcal{L}_{B\Box}$ is defined by the following grammar:*

$$\mathcal{L}_{B\Box} \ni \varphi ::= \mathsf{f} \mid \mathsf{t} \mid \top \mid \bot \mid p \in \mathsf{At} \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \supset \varphi \mid {\sim}\varphi \mid \Box\varphi,$$

The formulas of the language $\mathcal{L}_{B\Box}$ are interpreted in *four-valued Kripke models.*

Definition 4.2.2 (Four-valued Kripke frame (model)) *A four-valued Kripke frame $\mathcal{F} = \langle S, R\rangle$ is a pair, where S is a set of states and $R\colon S \times S \to$ FOUR is a* four-valued accessibility relation. *A four-valued Kripke model $\mathcal{M} = \langle \mathcal{F}, V\rangle$ is a pair, where $\mathcal{F} = \langle S, R\rangle$ is a four-valued Kripke frame and $V\colon \mathsf{At} \times S \to$ FOUR is a* four-valued valuation*, that assigns to each atomic proposition $p \in$ At and each state $s \in S$ a truth value from* FOUR. ⊣

The collection of four-valued Kripke models (frames) is denoted by fourMdl (fourFrm).

Four-valued Kripke models were defined in [98], but no notion of morphism was given there. We now introduce such a notion.

Definition 4.2.3 (fourMdl-bounded morphism) *Given two four-valued Kripke models $\mathcal{M}_1 = \langle S_1, R_1, V_1\rangle$ and $\mathcal{M}_2 = (S_2, R_2, V_2)$, a mapping $f\colon S_1 \to S_2$ is a* fourMdl-bounded morphism *(notation: $f\colon \mathcal{M}_1 \to \mathcal{M}_2$) if it satisfies the following conditions:*

1. *$V_1(p, s) = V_2(p, f(s))$, for every $p \in$ At and for every $s \in S$,*
2. *$R_1(s, t) = R_2(f(s), f(t))$, for all $s, t \in S$, and*
3. *for all $s \in S_1$ and for all $t \in S_2$, $R_2(f(s), t) = R_1(s, s')$, for some $s' \in S_1$ with $f(s') = t$.*

The definition of a fourFrm*-bounded morphism between four-valued Kripke frames is obtained by deleting the clause concerning valuations (item 1).* ⊣

Definition 4.2.4 (Extended four-valued valuation) *Let $\mathcal{M} = \langle S, R, V \rangle$ be a four-valued Kripke model. The extension of V to $\mathcal{L}_{B\Box}$ is a function $[\![-,-]\!]_{\mathcal{M}} \colon \mathcal{L}_{B\Box} \times S \to \mathsf{FOUR}$ that is inductively defined as follows:*

$$
\begin{aligned}
[\![\mathsf{t}, s]\!]_{\mathcal{M}} &= \mathsf{t} \\
[\![\top, s]\!]_{\mathcal{M}} &= \top \\
[\![\bot, s]\!]_{\mathcal{M}} &= \bot \\
[\![\mathsf{f}, s]\!]_{\mathcal{M}} &= \mathsf{f} \\
[\![p, s]\!]_{\mathcal{M}} &= V(p, s) \\
[\![\varphi \wedge \psi, s]\!]_{\mathcal{M}} &= [\![\varphi, s]\!]_{\mathcal{M}} \wedge [\![\psi, s]\!]_{\mathcal{M}} \\
[\![\varphi \vee \psi, s]\!]_{\mathcal{M}} &= [\![\varphi, s]\!]_{\mathcal{M}} \vee [\![\psi, s]\!]_{\mathcal{M}} \\
[\![\varphi \supset \psi, s]\!]_{\mathcal{M}} &= [\![\varphi, s]\!]_{\mathcal{M}} \supset [\![\psi, s]\!]_{\mathcal{M}} \\
[\![{\sim}\varphi, s]\!]_{\mathcal{M}} &= {\sim}[\![\varphi, s]\!]_{\mathcal{M}} \\
[\![\Box\varphi, s]\!]_{\mathcal{M}} &= \bigwedge_{t \in S} (R(s,t) \to [\![\varphi, t]\!]_{\mathcal{M}})
\end{aligned}
$$

where $\bigwedge$ denotes the infinitary version of $\wedge$ in FOUR *and $\to$ is the strong implication introduced in Table 4.1. We adopt the standard notational abbreviations for bilattice operations (Table 4.1). We employ $\Diamond\varphi$ as an abbreviation for ${\sim}\Box{\sim}\varphi$.* ⊣

For all four-valued Kripke models $\mathcal{M} = \langle S, R, V \rangle$ and state $s \in S$, it holds that

$$[\![\Diamond\varphi, s]\!]_{\mathcal{M}} = [\![{\sim}\Box{\sim}\varphi, s]\!]_{\mathcal{M}} = \bigvee_{t \in S} (R(s,t) * [\![\varphi, t]\!]_{\mathcal{M}}),$$

where $\bigvee$ denotes the infinitary version of $\vee$ in FOUR and $*$ is the strong conjunction. This shows that the two modal operators are inter-definable as in the classical case.

Definition 4.2.5 (Relational semantics for $\mathcal{L}_{B\Box}$) *Given a four-valued Kripke model $\mathcal{M} = \langle S, R, V \rangle$, a state $s \in S$, and a formula $\varphi \in \mathcal{L}_{B\Box}$, we say $(\mathcal{M}, s)$ satisfies φ and write $(\mathcal{M}, s) \models \varphi$ if $[\![\varphi, s]\!]_{\mathcal{M}} \in \{\mathsf{t}, \top\}$. This can be inductively defined as follows:*

$(\mathcal{M}, s) \models c$	*iff*	$c \in \{\mathsf{t}, \top\}$
$(\mathcal{M}, s) \models p$	*iff*	$V(p, s) \in \{\mathsf{t}, \top\}$
$(\mathcal{M}, s) \models \varphi \wedge \psi$	*iff*	$(\mathcal{M}, s) \models \varphi$ *and* $(\mathcal{M}, s) \models \psi$
$(\mathcal{M}, s) \models \varphi \vee \psi$	*iff*	$(\mathcal{M}, s) \models \varphi$ *or* $(\mathcal{M}, s) \models \psi$
$(\mathcal{M}, s) \models {\sim}\varphi$	*iff*	$(\mathcal{M}, s) \not\models \varphi$
$(\mathcal{M}, s) \models \varphi \supset \psi$	*iff*	$(\mathcal{M}, s) \models \varphi$ *implies* $(\mathcal{M}, s) \models \psi$
$(\mathcal{M}, s) \models \Box\varphi$	*iff*	*for all* $t \in S : R(s,t) \in \{\mathsf{t}, \top\}$ *implies* $(\mathcal{M}, t) \models \varphi$
$(\mathcal{M}, s) \models \Diamond\varphi$	*iff*	*there exists* $t \in S : R(s,t) \in \{\mathsf{t}, \top\}$ *and* $(\mathcal{M}, s) \models \varphi$.

A semantic consequence relation can now be introduced in the usual way. For a set of formulas $\Gamma \subseteq \mathcal{L}_{B\Box}$, we write $(\mathcal{M}, s) \models \Gamma$ to mean that $(\mathcal{M}, s) \models \gamma$ for each $\gamma \in \Gamma$. The (local) consequence $\Gamma \models_{\mathbf{BML}} \varphi$ holds if, for every model $\mathcal{M} = \langle S, R, V \rangle$ and every $s \in S$, it is the case that $(\mathcal{M}, s) \models \Gamma$ implies $(\mathcal{M}, s) \models \varphi$. For $\emptyset \models_{\mathbf{BML}} \varphi$ we write $\models_{\mathbf{BML}} \varphi$ (for 'φ is valid'). *The bilattice modal*

($\Box$**t**)	$\Box \mathsf{t} \leftrightarrow \mathsf{t}$
($\Box\wedge$)	$\Box(\varphi \wedge \psi) \leftrightarrow (\Box\varphi \wedge \Box\psi)$
($\Box\bot$)	$\Box(\bot \to \varphi) \leftrightarrow (\bot \to \Box\varphi)$
($\Box$−**monotonicity**)	from $\varphi \to \psi$ infer $\Box\varphi \to \Box\psi$

Table 4.4. The proof system $\mathbb{BML}$ consists of all axioms and rules of $\mathbb{LB}$ (Table 4.3) plus these three axioms and rule [98].

logic (**BML**) is defined as the set of valid formulas $\varphi \in \mathcal{L}_{B\Box}$. The consequence relation $\models_{\mathbf{BML}}$ inherits from the non-modal fragment the deduction theorem in the following form: $\Gamma \models_{\mathbf{BML}} \varphi$ if and only if there is a finite $\Gamma' \subseteq \Gamma$ such that $\models_{\mathbf{BML}} \bigwedge \Gamma' \supset \varphi$, where $\bigwedge \Gamma' := \bigwedge\{\gamma \in \Gamma'\}$. It implies that in the axiomatisation task, one can without loss of generality restrict attention to valid formulas. This consequence relation is axiomatised in [98]. The axiomatisation $\mathbb{BML}$ is displayed in Table 4.4.

A derivation in the proof system $\mathbb{BML}$ is a sequence of formulas such that every formula is an instantiation of an axiom or the result of applying a rule to formulas prior in the sequence. If φ occurs in a derivation we write $\vdash_{\mathbb{BML}} \varphi$, for "$\varphi$ is a theorem". By $\Gamma \vdash_{\mathbb{BML}} \varphi$ we mean that there is a finite subset Γ' of Γ such that $\vdash_{\mathbb{BML}} \bigwedge \Gamma' \supset \varphi$.

We note that the *necessitation rule* "from φ, infer $\Box\varphi$" is not valid [98, Section III.A], and the *normality axiom* $\Box(\varphi \supset \psi) \supset (\Box\varphi \supset \Box\psi)$ also is not valid [28]. In [98], the following soundness and completeness result with respect to the four-valued Kripke semantics was shown.

Theorem 4.2.6 (Relational soundness and completeness [98, Theorem 19]) *For all* $\Gamma, \{\varphi\} \subseteq \mathcal{L}_{B\Box}$, $\Gamma \vdash_{\mathbb{BML}} \varphi$ *iff* $\Gamma \models_{\mathbf{BML}} \varphi$. ⊣

4.2.3 Bilattice modal logic: algebraic semantics

We recall an algebraic semantics for bilattice modal logic based on *modal bilattices* [98]. We give a brief account of the algebraic completeness for $\mathbb{BML}$, as we will build on it later on. Let us first recall some basic definitions.

Definition 4.2.7 (Bilattice) *A* (bounded) distributive bilattice *is an algebra* $\mathbf{B} = \langle B, \wedge, \vee, \otimes, \oplus, \sim, \mathsf{t}, \top, \bot, \mathsf{f}\rangle$, *where* $\langle B, \wedge, \vee, \mathsf{t}, \mathsf{f}\rangle$ *and* $\langle B, \otimes, \oplus, \top, \bot\rangle$ *are (bounded) distributive lattices [44], and for all* $x, y, z \in B$*, the following identity is satisfied:*

$$x \circ (y \bullet z) = (x \circ z) \bullet (x \circ y) \qquad \circ, \bullet \in \{\wedge, \vee, \otimes, \oplus\}.$$

The bilattice negation $\sim$ *is required to satisfy the following conditions: for every* $x, y \in B$

1. $x \leq_t y \iff {\sim}y \leq_t {\sim}x$,
2. $x \leq_k y \iff {\sim}x \leq_k {\sim}y$, *and*
3. ${\sim}{\sim}x = x$. ⊣

Similar to Belnap's lattice FOUR*, the order* $\leq_t$ *arising from* $\wedge$ *and* $\vee$ *is called the truth-order, and the order* $\leq_k$ *arising from* $\otimes$ *and* $\oplus$ *is called the knowledge order.*

The conditions (1)-(3) determine the behaviour of the bilattice negation on FOUR: ${\sim}\mathsf{t} = \mathsf{f}$, ${\sim}\mathsf{f} = \mathsf{t}$, ${\sim}\top = \top$, and ${\sim}\bot = \bot$. Hence, FOUR is the smallest non-trivial bilattice. Figure 4.2 depicts some examples of bilattices. We may label the operations and constants with the name of the bilattice, as in $\Box_{\mathbf{B}}$, $\mathsf{f}_{\mathbf{B}}$, etc., to distinguish them from those in other bilattices.

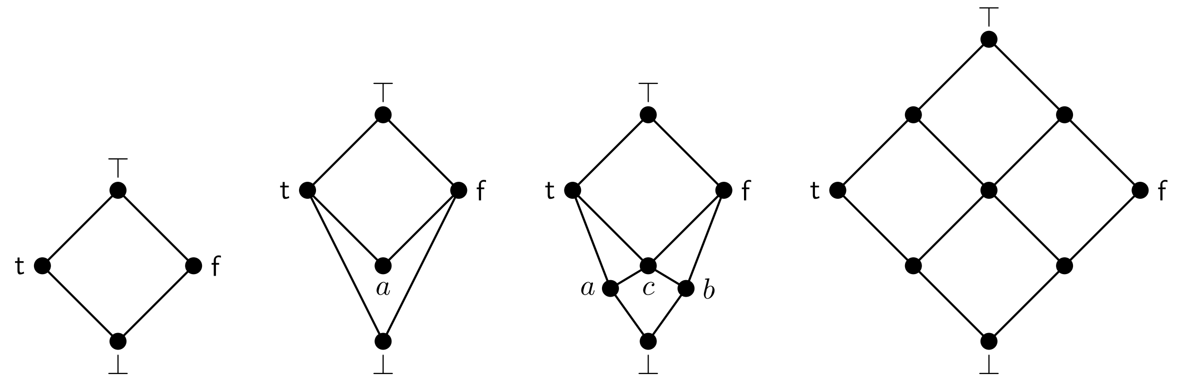

(a) FOUR (b) FIVE (c) SEVEN (d) NINE

Figure 4.2. Some examples of bilattices

Definition 4.2.8 (Bilattice homomorphism) *Given two bilattices* $\mathbf{B}$ *and* $\mathbf{B}'$*, a function* $f\colon B \to B'$ *is a* bilattice homomorphism *(notation:* $f\colon \mathbf{B} \to \mathbf{B}'$*), if* f *is a lattice homomorphism with respect to both lattices* $\langle B, \wedge, \vee, \sim, \mathsf{t}, \mathsf{f}\rangle$ *and* $\langle B, \otimes, \oplus, \sim, \top, \bot\rangle$*.* ⊣

The collection of bilattices is denoted by Bilat.

Definition 4.2.9 (Implicative bilattice) *An implicative bilattice is an algebraic structure* $\mathbf{B} = \langle B, \wedge, \vee, \otimes, \oplus, \sim, \supset, \mathsf{t}, \top, \bot, \mathsf{f}\rangle$*, where* $\langle B, \wedge, \vee, \otimes, \oplus, \sim, \mathsf{t}, \top, \bot, \mathsf{f}\rangle$ *is a bilattice, and the operation* $\supset$ *satisfies the following identities: for every* $x, y, z \in B$

1. $(x \supset x) \supset y = y,$
2. $x \supset (y \supset z) = (x \wedge y) \supset z = (x \oplus y) \supset z,$
3. $((x \supset y) \supset x) \supset x = x \supset x,$
4. $(x \vee y) \supset z = (x \supset z) \wedge (y \supset z) = (x \oplus y) \supset z,$
5. $x \supset (y \supset z) = (x \wedge y) \supset z = (x \otimes y) \supset z,$
6. $x \wedge ((x \supset y) \supset (x \oplus y)) = x,$
7. $\sim(x \supset y) \supset z = (x \wedge \sim y) \supset z.$ ⊣

It is easy to check that FOUR, viewed as an algebra in the language $\langle \wedge, \vee, \otimes, \oplus, \sim, \supset, \mathsf{t}, \top, \bot, \mathsf{f}\rangle$, is the smallest non-trivial implicative bilattice.

Let $\mathbf{B}$ and $\mathbf{B}'$ be two implicative bilattices. A mapping $f\colon B \to B'$ is a *imp-bilattice homomorphism*, if it is a bilattice homomorphism and $f(x \supset^{\mathbf{B}} y) = f(x) \supset^{\mathbf{B}'} f(y)$, for every $x, y \in B$. The collection of implicative bilattices is denoted by ImpBilat. Now, we define *modal bilattices*.

Definition 4.2.10 (Modal bilattice) *A* modal bilattice *is an algebra*

$$\mathbb{B} = \langle B, \wedge, \vee, \otimes, \oplus, \supset, \sim, \mathsf{t}, \top, \bot, \mathsf{f}, \Box\rangle$$

where the $\Box$*-free reduct of* $\mathbb{B}$ *is an implicative bilattice and the following identities are satisfied:*

1. $\Box \mathsf{t} = \mathsf{t}$*,*

2. $\Box(x \wedge y) = \Box x \wedge \Box y$*, and*

3. $\Box(\bot \to x) = \bot \to \Box x$*.* ⊣

We note that the identities (1)-(3) correspond, respectively, to axioms $\Box\mathbf{t}$, $\Box\wedge$ and $\Box\bot$ of the calculus $\mathbb{BML}$ (Table 4.4). For notational convenience, we sometimes write $\mathbb{B} = \langle \mathbf{B}, \Box \rangle$ instead of $\mathbb{B} = \langle B, \wedge, \vee, \otimes, \oplus, \supset, \sim, \mathsf{t}, \top, \bot, \mathsf{f}, \Box \rangle$, where $\mathbf{B}$ is the $\Box$-free reduct of $\mathbb{B}$, i.e., $\mathbf{B} = \langle B, \wedge, \vee, \otimes, \oplus, \supset, \sim, \mathsf{t}, \top, \bot, \mathsf{f} \rangle$.

Definition 4.2.11 (Modal bilattice homomorphism) *Given two modal bilattices* $\mathbb{B} = \langle \mathbf{B}, \Box_{\mathbb{B}} \rangle$ *and* $\mathbb{B}' = \langle \mathbf{B}', \Box_{\mathbb{B}'} \rangle$*, a mapping* $f : B \to B'$ *is a* modal bilattice homomorphism *(notation* $f \colon \mathbb{B} \to \mathbb{B}'$*), if* f *is an imp-bilattice homomorphism between* $\mathbf{B}$ *and* $\mathbf{B}'$*, and* $f(\Box_{\mathbb{B}} x) = \Box_{\mathbb{B}'} f(x)$*, for every* $x \in B$*.* ⊣

The collection of modal bilattices is denoted by MBilat.

Definition 4.2.12 (Bifilter) *A subset* $F \subseteq B$ *of a modal bilattice* $\mathbb{B}$ *is a* bifilter*, if* F *is a lattice filter with respect to both orders* $\leq_t$ *and* $\leq_k$ *[29, Prop. 2.11].* ⊣

Given a pair $\langle \mathbb{B}, F \rangle$, where $\mathbb{B}$ is a modal bilattice and F is a bifilter of $\mathbb{B}$, and sets of formulas $\Gamma, \{\varphi\} \subseteq \mathcal{L}_{B\Box}$, we write $\Gamma \models_{\langle B,F \rangle} \varphi$ to mean that, for all functions $h \colon \mathcal{L}_{B\Box} \to \mathbb{B}$, if $h(\gamma) \in F$ for all $\gamma \in \Gamma$, then also $h(\varphi) \in F$, where h is a *logical homomorphism*, i.e., for all $\varphi, \psi \in \mathcal{L}_{B\Box}$,

$$\begin{aligned} h(c) &= c_{\mathbb{B}} && (c \in \{\mathsf{t}, \top, \bot, \mathsf{f}\}) \\ h(\sim\varphi) &= \sim_{\mathbb{B}} h(\varphi) \\ h(\Box\varphi) &= \Box_{\mathbb{B}} h(\varphi) \\ h(\varphi \circ \psi) &= h(\varphi) \circ_{\mathbb{B}} h(\psi). && (\circ \in \{\wedge, \vee, \otimes, \oplus, \supset\}) \end{aligned}$$

A formula $\varphi \in \mathcal{L}_{B\Box}$ is valid over the class of modal bilattices, if for all modal bilattices $\mathbb{B}$ and for all logical homomorphisms $h \colon \mathcal{L}_{B\Box} \to B$, $h(\varphi) \geq_t \top_{\mathbb{B}}$.

We can then state algebraic soundness and completeness [98, Theorem 10].

Theorem 4.2.13 (Algebraic soundness and completeness) *For all* $\Gamma, \{\varphi\} \subseteq \mathcal{L}_{B\Box}$*,* $\Gamma \vdash_{\mathbb{BML}} \varphi$ *iff for all modal bilattices* $\mathbb{B}$ *and all bifilters* $F \subseteq B$*,* $\Gamma \models_{\langle \mathbb{B},F \rangle} \varphi$ *.* ⊣

4.2.4 Duality for modal bilattices

Just as with basic modal logic, the relational and the algebraic semantics for bilattice modal logic are related via a Stone-type duality [98, Theorem 18]. In the case of bilattices, another key ingredient that greatly simplifies the picture is the so-called *twist structure* representation, which works as follows. We first recall *bimodal Boolean algebras.*

Definition 4.2.14 (Bimodal Boolean algebra) *An algebra* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1} \rangle$ *is a* bimodal Boolean algebra*, if* $\langle A, \wedge, \vee, \neg, \Box^+, \mathbf{0}, \mathbf{1} \rangle$ *and* $\langle A, \wedge, \vee, \neg, \Box^-, \mathbf{0}, \mathbf{1} \rangle$ *are both modal Boolean algebras (Def. 3.3.1). Note that no relation between* $\Box^+$ *and* $\Box^-$ *is assumed.* ⊣

To simplify the notation, we will sometimes write $\langle \mathbf{A}, \Box^+, \Box^- \rangle$ instead of $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1} \rangle$, where $\mathbf{A} = \langle A, \wedge, \vee, \neg, \mathbf{0}, \mathbf{1} \rangle$ is a Boolean algebra. The dual operations $\Diamond^+$ and $\Diamond^-$ are defined in the usual way by setting $\Diamond^+ x := \neg\Box^+\neg x$ and $\Diamond^- x := \neg\Box^-\neg x$. We denote by $\mathsf{2MBA}$ the collection of bimodal Boolean algebras.

Definition 4.2.15 (2MBA-homomorphism) *Given two bimodal Boolean algebras* $\mathbb{A}_1 = \langle \mathbf{A}_1, \Box_1^+, \Box_1^- \rangle$ *and* $\mathbb{A}_2 = \langle \mathbf{A}_2, \Box_2^+, \Box_2^- \rangle$*, a function* $f \colon A_1 \to A_2$ *is a* 2MBA-homomorphism *(notation:* $f \colon \mathbb{A}_1 \to \mathbb{A}_2$*), if it is a modal Boolean algebra homomorphism from* $\langle \mathbf{A}_1, \Box_1^+ \rangle$ *to* $\langle \mathbf{A}_2, \Box_2^+ \rangle$ *and from* $\langle \mathbf{A}_1, \Box_1^- \rangle$ *to* $\langle \mathbf{A}_2, \Box_2^- \rangle$*.* ⊣

Definition 4.2.16 (Twist structure) *Let* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1} \rangle$ *be a bimodal Boolean algebra. The* twist structure over $\mathbb{A}$ *is defined as the algebra*

$$\mathbb{A}^{\bowtie} = \langle A \times A, \wedge, \vee, \otimes, \oplus, \supset, \sim, \Box, \mathsf{t}, \top, \bot, \mathsf{f} \rangle$$

with operations given, for all $\langle x_1, x_2 \rangle, \langle y_1, y_2 \rangle \in A \times A$*, by:*

$$\begin{aligned}
\langle x_1, x_2 \rangle \wedge \langle y_1, y_2 \rangle &:= \langle x_1 \wedge y_1,\ x_2 \vee y_2 \rangle \\
\langle x_1, x_2 \rangle \vee \langle y_1, y_2 \rangle &:= \langle x_1 \vee y_1,\ x_2 \wedge y_2 \rangle \\
\langle x_1, x_2 \rangle \otimes \langle y_1, y_2 \rangle &:= \langle x_1 \wedge y_1,\ x_2 \wedge y_2 \rangle \\
\langle x_1, x_2 \rangle \oplus \langle y_1, y_2 \rangle &:= \langle x_1 \vee y_1,\ x_2 \vee y_2 \rangle \\
\langle x_1, x_2 \rangle \supset \langle y_1, y_2 \rangle &:= \langle \neg x_1 \vee y_1, x_1 \wedge y_2 \rangle \\
\sim\langle x_1, x_2 \rangle &:= \langle x_2,\ x_1 \rangle \\
\Box\langle x_1, x_2 \rangle &:= \langle \Box^+ x_1 \wedge \Box^- \neg x_2,\ \Diamond^+ x_2 \rangle \\
\mathsf{t} &:= \langle \mathbf{1}, \mathbf{0} \rangle \\
\top &:= \langle \mathbf{1}, \mathbf{1} \rangle \\
\bot &:= \langle \mathbf{0}, \mathbf{0} \rangle \\
\mathsf{f} &:= \langle \mathbf{0}, \mathbf{1} \rangle
\end{aligned}$$

Every twist structure $\mathbb{A}^{\bowtie}$ over a bimodal Boolean algebra $\mathbb{A}$ is a modal bilattice [140, Prop. 5.17], in which the dual modality of $\Box$ is defined as follows:

$$\Diamond\langle x_1, x_2 \rangle := \langle \Diamond^+ x_1,\ \Box^+ x_2 \wedge \Box^- \neg x_1 \rangle.$$

Conversely, any modal bilattice $\mathbb{B}$ is isomorphic to a twist structure of a bimodal Boolean algebra $\mathbb{A}$ [140, Theorem 5.18].

Definition 4.2.17 *Let* $\mathbb{B} = \langle B, \wedge, \vee, \otimes, \oplus, \sim, \supset, \mathsf{t}, \top, \bot, \mathsf{f}, \Box \rangle$ *be a modal bilattice. We construct a bimodal Boolean algebra* $\mathbb{B}_{\bowtie}$ *by defining an equivalence relation over* B *as follows: for all* $a, b \in B$

$$a \equiv b \quad \iff \quad a \wedge b = a \oplus b. \tag{4.1}$$

The set of equivalence classes $B/\equiv$ *can be endowed with the following operators: for all* $[a], [b] \in B/\equiv$

$$\begin{aligned}
[a] \wedge [b] &= [a \wedge b] \\
[a] \vee [b] &= [a \vee b] \\
\neg[a] &= [a \supset \mathsf{f}] \\
\Box^+[a] &= [\Diamond(a \supset \mathsf{f}) \supset \mathsf{f}] \\
\Box^-[a] &= [\Box(\sim(a \supset \mathsf{f}) \vee \top)] \\
\mathbf{1} &= [\mathsf{t}] \\
\mathbf{0} &= [\mathsf{f}].
\end{aligned}$$

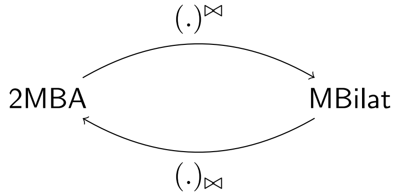

Figure 4.3. Transformations between bimodal Boolean algebras and modal bilattices

It takes a simple check to verify that the algebra

$$\mathbb{B}_{\bowtie} = \langle B/\equiv, \wedge, \vee, \neg, \Box^{+}, \Box^{-}, \mathbf{1}, \mathbf{0} \rangle \tag{4.2}$$

is a bimodal Boolean algebra. ⊣

Then, [140, 5.18] shows that for every modal bilattice $\mathbb{B}$ and every bimodal Boolean algebra $\mathbb{A}$, we have

$$\mathbb{B} \cong (\mathbb{B}_{\bowtie})^{\bowtie} \quad \text{and} \quad \mathbb{A} \cong (\mathbb{A}^{\bowtie})_{\bowtie}. \tag{4.3}$$

through the isomorphism $j^{\bowtie} : \mathbb{B} \to (\mathbb{B}_{\bowtie})^{\bowtie}$ which maps every $a \in B$ to the pair $([a], [\sim a])$. The relation between the MBilat and 2MBA is depicted in Figure 4.3.

Remark 4.2.18 Readers familiar with category theory may verify that the operation $(.)^{\bowtie}$ is a functor from the category of bimodal Boolean algebras with 2MBA-homomorphisms to the category of modal bilattices with modal bilattice homomorphisms, and vice versa for $(.)_{\bowtie}$. Moreover,it can be shown that these functors, in fact, form an equivalence [98]. This categorical perspective is implicit in what follows, but seldom comes to the surface. ⊣

The twist structure construction allows us to relate four-valued Kripke frames and modal bilattices via Jónsson-Tarski duality for basic modal logic (see, e.g., [78]). We first recall relevant definitions.

Definition 4.2.19 (Bimodal Kripke frame (model)) *A* bimodal Kripke frame *is a tuple* $\mathcal{F} = \langle S, R^{+}, R^{-} \rangle$*, where* $\langle S, R^{+} \rangle$ *and* $\langle S, R^{-} \rangle$ *are Kripke frames. Similarly, a* bimodal Kripke model *is a tuple* $\mathcal{M} = \langle \mathcal{F}, V^{+}, V^{-} \rangle$*, where* $\mathcal{F} = \langle S, R^{+}, R^{-} \rangle$ *is a bimodal Kripke frame and* $V^{+}, V^{-} \colon \mathsf{At} \to \mathcal{P}(S)$ *are valuations. The collection of bimodal Kripke frames (models) is denoted by* 2Frm *(*2Mdl*).*

Definition 4.2.20 (2Mdl-bounded morphism) *Given two bimodal Kripke models* $\mathcal{M}_1 = \langle S_1, R_1^{+}, R_1^{-}, V_1^{+}, V_1^{-} \rangle$ *and* $\mathcal{M}_2 = \langle S_2, R_2^{+}, R_2^{-}, V_2^{+}, V_2^{+} \rangle$*, a mapping* $f \colon S_1 \to S_2$ *is a* 2Mdl-bounded morphism *(notation:* $f : \mathcal{M}_1 \to \mathcal{M}_2$*), if it is a bounded morphism [25, Def. 2.10] between* $\langle S_1, R_1^{+}, V_1^{+} \rangle$ *and* $\langle S_2, R_2^{+}, V_2^{+} \rangle$*, and between* $\langle S_1, R_1^{-}, V_1^{-} \rangle$ *and* $\langle S_2, R_2^{-}, V_2^{-} \rangle$*. Similarly, a* 2Frm-*bounded morphism between two bimodal Kripke frames* $\mathcal{F}_1 = \langle S_1, R_1^{+}, R_1^{-} \rangle$ *and* $\mathcal{F}_2 = \langle S_2, R_2^{+}, R_2^{-} \rangle$ *(notation:* $f : \mathcal{F}_1 \to \mathcal{F}_2$*) is a bounded morphism between* $\langle S_1, R_1^{+} \rangle$ *and* $\langle S_2, R_2^{+} \rangle$ *and between* $\langle S_1, R_1^{-} \rangle$ *and* $\langle S_2, R_1^{-} \rangle$*.* ⊣

It is shown in [138, 139] that there is a correspondence between fourFrm (fourMdl) and 2Frm (2Mdl).

Definition 4.2.21 *Given a four-valued Kripke frame $\mathcal{F} = \langle S, R\rangle$, we define the bimodal Kripke frame*

$$\mathcal{F}_{\bowtie} = (S, R^+, R^-)$$

where $R^+, R^- \subseteq S \times S$ are defined as follows: for every $s, t \in S$,

$$t \in R^+(s) \iff R(s,t) \in \{\mathsf{t}, \top\} \quad \text{and} \quad t \in R^-(s) \iff R(s,t) \in \{\mathsf{t}, \bot\}. \tag{4.4}$$

Given a four-valued Kripke model $\mathcal{M} = \langle \mathcal{F}, V\rangle$, we define the bimodal Kripke model

$$\mathcal{M}_{\bowtie} = \langle \mathcal{F}_{\bowtie}, V^+, V^-\rangle$$

where $V^+, V^- \colon \mathsf{At} \to \mathcal{P}(S)$ are defined as follows: for every $p \in \mathsf{At}$,

$$V^+(p) = \{s \in S : V(p,s) \in \{\mathsf{t}, \top\}\} \quad \text{and} \quad V^-(p) = \{s \in S : V(p,s) \in \{\mathsf{f}, \top\}\}. \tag{4.5}$$

We note that $V^+(p) = V^-(\sim p)$, for every $p \in \mathsf{At}$. ⊣

Definition 4.2.22 *Given a bimodal Kripke frame $\mathcal{F} = \langle S, R^+, R^-\rangle$, we define the four-valued Kripke frame $\mathcal{F}^{\bowtie} = \langle S, R\rangle$ where $R \colon S \times S \to \mathsf{FOUR}$ is a four-valued relation defined by:*

$$R(s,t) = \begin{cases} \mathsf{t} & \text{if } (s,t) \in R^+ \cap R^-, \\ \top & \text{if } (s,t) \in R^+ \setminus R^-, \\ \bot & \text{if } (s,t) \in R^- \setminus R^+, \\ \mathsf{f} & \text{if } (s,t) \notin R^+ \cup R^-. \end{cases}$$

Similarly, for every bimodal Kripke model $\mathcal{M} = \langle \mathcal{F}, V^+, V^-\rangle$ we define the four-valued Kripke model $\mathcal{M}^{\bowtie} = \langle \mathcal{F}^{\bowtie}, V\rangle$ where $V \colon \mathsf{At} \times S \to \mathsf{FOUR}$ is a four-valued valuation defined by: for all $s, t \in S$ and $p \in \mathsf{At}$

$$V(p,s) = \begin{cases} \mathsf{t} & \text{if } s \in V^+(p) \setminus V^-(p), \\ \top & \text{if } s \in V^+(p) \cap V^-(p), \\ \bot & \text{if } s \notin V^+(p) \cup V^-(p), \\ \mathsf{f} & \text{if } s \in V^-(p) \setminus V^+(p). \end{cases}$$

Proposition 4.2.23 *Let $\mathcal{M}_1 = \langle S_1, R_1, V_1\rangle$ and $\mathcal{M}_2 = \langle S_2, R_2, V_2\rangle$ be two four-valued Kripke models. Every $\mathsf{fourMdl}$-bounded morphism $f : \mathcal{M}_1 \to \mathcal{M}_2$ is a $\mathsf{2Mdl}$-bounded morphism from $(\mathcal{M}_1)_{\bowtie} = \langle S_1, R_1^+, R_1^-, V_1^+, V_1^-\rangle$ to $(\mathcal{M}_2)_{\bowtie} = \langle S_2, R_2^+, R_2^-, V_2^+, V_2^-\rangle$.* ⊣

Proof Let $\mathcal{M}_1 = \langle S_1, R_1, V_1\rangle$ and $\mathcal{M}_2 = \langle S_2, R_2, V_2\rangle$ be two four-valued Kripke models, and let $f : \mathcal{M}_1 \to \mathcal{M}_2$ be a bounded morphism. We need to show that for $\sigma \in \{+, -\}$, f is a bounded morphism from $\langle S_1^\sigma, R_1^\sigma, V_1^\sigma\rangle$ to $\langle S_2^\sigma, R_2^\sigma, V_2^\sigma\rangle$. We only show the case for $\sigma = +$. The case for $\sigma = -$ is similar. So we show that:

1. For all $s \in S_1$, $s \in V_1^+(p)$ iff $f(s) \in V_2^+(p)$.

2. For all $s, t \in S_1$: if $t \in R_1^+(s)$ then $f(t) \in R_2^+(f(s))$.

3. For all $s, t \in S_1$: if $t \in R_2^+(f(s))$ then there is an $s' \in S_1$ such that $s' \in R_1^+(s)$ and $f(s') = t$.

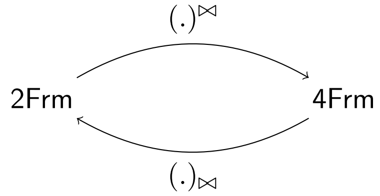

Figure 4.4. Transformation between four-valued Kripke frames and bimodal Kripke frames

Item 1 Let $s \in S_1$ and $p \in \mathsf{At}$. We have

$$\begin{aligned} s \in V_1^+(p) &\iff V_1(p,s) \in \{\mathsf{t}, \top\} \\ &\iff V_2(p, f(s)) \in \{\mathsf{t}, \top\} \quad \text{(Def. 4.2.3, } f \text{ is a } \mathsf{fourMdl}\text{-bounded morphism)} \\ &\iff f(s) \in V_2^+(p). \end{aligned}$$

Item 2 Let $s, t \in S_1$ be such that $t \in R_1^+(s)$. By the definition of R_1^+ (see (4.4)), it means that $R_1(s,t) \in \{\mathsf{t}, \top\}$. Since f is a four-valued bounded morphism, it follows that $R_2(f(s), f(t)) \in \{\mathsf{t}, \top\}$. Hence, by (4.4) we have that $f(t) \in R_2^+(f(s))$.

Item 3 Let $s \in S_1$ and $t \in S_2$ be such that $t \in R_2^+(f(s))$. By (4.4) we have that $R_2(f(s), t) \in \{\mathsf{t}, \top\}$. Since f is a four-valued bounded morphism, there exists $s' \in S_1$ such that $f(s') = t$ and $R_2(f(s), t) = R_1(s, s')$. It then follows from (4.4) that $s' \in R_1^+(s)$ and this completes the proof.

□

Proposition 4.2.24 *Let* $\mathcal{M}_1 = \langle S_1, R_1^+, R_1^-, V_1^+, V_1^- \rangle$ *to* $\mathcal{M}_2 = \langle S_2, R_2^+, R_2^-, V_2^+, V_2^- \rangle$ *be bimodal Kripke models, and let* $f : \mathcal{M}_1 \to \mathcal{M}_2$ *be a* 2Mdl*-bounded morphism. Then* f *is a* fourMdl*-bounded morphism from* $\mathcal{M}_1^{\bowtie} = \langle S_1, R_1, V_1 \rangle$ *to* $\mathcal{M}_2^{\bowtie} = \langle S_2, R_2, V_2 \rangle$. ⊣

Proof This proof is analogous to that of Prop. 4.2.23. □

One can easily show that for every four-valued Kripke frame $\mathcal{F}$, and four-valued Kripke model, we have

$$\mathcal{F} = (\mathcal{F}_{\bowtie})^{\bowtie} \quad \text{and} \quad \mathcal{M} = (\mathcal{M}_{\bowtie})^{\bowtie}. \tag{4.6}$$

Furthermore, for every bimodal Kripke frame $\mathcal{F}$ and bimodal Kripke model $\mathcal{M}$,

$$\mathcal{F} = (\mathcal{F}^{\bowtie})_{\bowtie} \quad \text{and} \quad \mathcal{M} = (\mathcal{M}^{\bowtie})_{\bowtie}. \tag{4.7}$$

The relation between bimodal Kripke frames and four-valued Kripke frames is depicted in Figure 4.4.

Remark 4.2.25 Again, the reader who is familiar with category theory will have noticed that (4.6) and (4.7) with Prop. 4.2.23 and Prop. 4.2.24 show that the category of four-valued Kripke models with fourMdl-bounded morphisms is isomorphic to the category of bimodal Kripke models with bimodal bounded morphisms. ⊣

Duality between four-valued Kripke frames and modal bilattices. The *complex duality* (or *discrete duality*) between so-called perfect modal Boolean algebras and Kripke frames (cf. [26, Section 5.2] and [152]) can be easily extended to bimodal Kripke frames and perfect bimodal Boolean algebras. To this end, we first recall *perfect modal Boolean algebras* and *perfect modal bilattices* from [139, Section 5.2].

Definition 4.2.26 (Perfect modal Boolean algebra) *A modal Boolean algebra* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box, \mathbf{0}, \mathbf{1}\rangle$ *is called* perfect *if (i)* $\mathbb{A}$ *is complete, (ii) atomic, i.e.,* $\mathbb{A}$ *is completely join-generated by its set of* atoms $\mathsf{At}(\mathbb{A}) := \{x \in A : x \neq 0,$ *and for all* $y \in A$, $y < x$ *implies* $y = 0\}$, *and* $\Box$ *preserves infinitary meets. The collection of perfect modal Boolean algebras is denoted by* PrMBA. ⊣

One can easily check that the complex algebra of a Kripke frame is a perfect modal Boolean algebra.

Definition 4.2.27 (Perfect bimodal Boolean algebra) *A bimodal Boolean algebra* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1}\rangle$ *is perfect if both* $\langle A, \wedge, \vee, \neg, \Box^+, \mathbf{0}, \mathbf{1}\rangle$ *and* $\langle A, \wedge, \vee, \neg, \Box^-, \mathbf{0}, \mathbf{1}\rangle$ *are perfect modal Boolean algebras. We denote by* $\mathsf{Pr2MBA}$ *the collection of perfect bimodal Boolean algebras.* ⊣

Using the twist structure representation of modal bilattices, we can define a perfect modal bilattice as follows.

Definition 4.2.28 (Perfect modal bilattice) *A modal bilattice* $\mathbb{B}$ *is called* perfect *if* $\mathbb{B} \cong \mathbb{A}^{\bowtie}$, *where* $\mathbb{A}$ *is a perfect bimodal Boolean algebra. The collection of perfect modal bilattices is denoted by* $\mathsf{PrMBilat}$. ⊣

In order to establish the duality between four-valued Kripke frames and modal bilattices, we first describe the duality between bimodal Kripke frames and perfect bimodal Boolean algebras which is easily obtained from the basic complex duality.

Definition 4.2.29 *Let* $\mathcal{F} = \langle S, R^+, R^-\rangle$ *be a bimodal Kripke frame. The* complex bimodal algebra *of* $\mathcal{F}$ *is the bimodal Boolean algebra*

$$\mathcal{F}^\bullet = \langle \mathcal{P}(S), \cap, \cup, (-)^c, \Box^+, \Box^-, \emptyset, S\rangle,$$

where $\mathcal{P}(S)$ *is the powerset of* S, *and* $\Box^+$ *and* $\Box^-$ *are defined as follows: for every* $X \subseteq S$

$$\Box^+ X = \{s \in S : R^+(s) \subseteq X\}, \quad \text{and} \quad \Box^- X = \{s \in S : R^-(s) \subseteq X\}. \tag{4.8}$$

Notice that $\langle \mathcal{P}(S), \cap, \cup, (-)^c, \Box^+, S, \emptyset\rangle$, $\langle \mathcal{P}(S), \cap, \cup, (-)^c, \Box^-, S, \emptyset\rangle$ *are the complex algebras of the two-valued Kripke frames* $\langle S, R^+\rangle$ *and* $\langle S, R^-\rangle$, *respectively.*

Going in the opposite direction, let $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1}\rangle$ *be a perfect bimodal Boolean algebra. The dual bimodal Kripke frame of* $\mathbb{A}$ *is defined as*

$$\mathbb{A}_\bullet = \langle \mathsf{At}(\mathbb{A}), R^+, R^-\rangle,$$

where R^+ *and* R^- *are defined by*

$$R^+(s) = \{s' \in \mathsf{At}(\mathbb{A}) : s \leq \Diamond^+ s'\}, \quad \text{and} \quad R^-(s) = \{s' \in \mathsf{At}(\mathbb{A}) : s \leq \Diamond^- s'\}. \tag{4.9}$$

where $\Diamond^+ s = \neg\Box^+\neg s$ *and* $\Diamond^- s = \neg\Box^-\neg s$, *for every* $s \in \mathsf{At}(\mathbb{A})$. ⊣

Let $\mathcal{F}$ be a bimodal Kripke frame, and let $\mathbb{A}$ be a perfect bimodal Boolean algebra. Then, it follows from basic complex duality theory, see e.g., [26, Section 5.2] and [152] that

$$(\mathbb{A}_\bullet)^\bullet \cong \mathbb{A} \quad \text{and} \quad (\mathcal{F}^\bullet)_\bullet \cong \mathcal{F}. \tag{4.10}$$

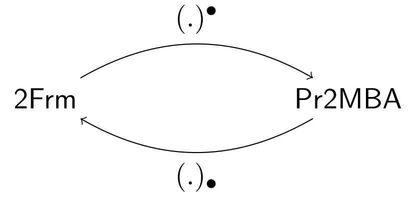

Figure 4.5. Transformations between bimodal Kripke frames and perfect bimodal Boolean algebras.

Remark 4.2.30 The above correspondence between bimodal Kripke frames and perfect bimodal Boolean algebras is, in fact, just the object part of the dual equivalence between the category of bimodal Kripke frames with bimodal bounded morphisms and the category of bimodal Boolean algebras with 2MBA-homomorphisms. ⊣

By composing the constructions we have described so far, we obtain constructions between four-valued Kripke frames and perfect modal bilattices. The constructions are summarised in Figure 4.6.

Definition 4.2.31 *Let $\mathbb{B}$ be a modal bilattice. We define its dual four-valued Kripke frame as*

$$\mathbb{B}_* = ((\mathbb{B}_{\bowtie})_\bullet)^{\bowtie}.$$

On the other hand, for every four-valued Kripke frame $\mathcal{F}$, we define its complex bilattice *by*

$$\mathcal{F}^* = ((\mathcal{F}_{\bowtie})^\bullet)^{\bowtie}.$$

Moreover, for every four-valued Kripke model $\mathcal{M} = \langle \mathcal{F}, V \rangle$, where $V \colon \mathsf{At} \times S \to \mathsf{FOUR}$ is a four-valued valuation, we define

$$\mathcal{M}^* = \langle \mathcal{F}^*, V^* \rangle$$

where $\mathcal{F}^*$ is the complex bilattice of $\mathcal{F}$ and $V^* \colon \mathsf{At} \to \mathcal{F}^*$ is defined by

$$V^*(p) = \langle V^+(p), V^-(p) \rangle$$

where $V^+, V^- \colon \mathsf{At} \to \mathcal{P}(S)$ are the valuations that are defined in (4.5). ⊣

We note that since complex bimodal algebras are perfect, it follows that complex bilattices are perfect. The following theorem summarises the object part of the duality between perfect modal bilattices and four-valued Kripke frames [139, Proposition 5.5].

Theorem 4.2.32 *For every four-valued Kripke frame $\mathcal{F}$ and every perfect modal bilattice $\mathbb{B}$, we have that*

$$\mathcal{F} \cong (\mathcal{F}^*)_* \quad \textit{and} \quad \mathbb{B} \cong (\mathbb{B}_*)^* \tag{4.11}$$

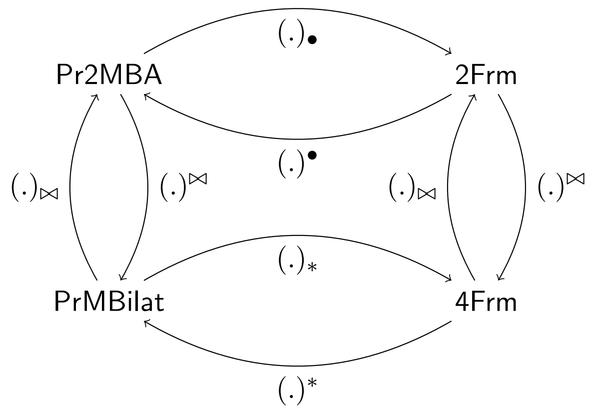

Figure 4.6. Transformations between four-valued Kripke frames and perfect modal bilattices.

4.3 The bilattice action model logic: syntax and semantics

In this section we introduce the bilattice action model logic (**BAML**). In order to obtain an algebraic semantics, we apply the methods of [114, 105] that we have explained in Section 3.4 (see page 37) to the class of modal bilattices using insights of [138, 139] and the relationship between modal bilattices and twist structures. By doing so, we obtain the main ingredient of the algebraic semantics of **BAML**, the notion of an *intermediate bilattice*. We then use the duality between four-valued Kripke models and perfect modal bilattices to define a relational semantics based on four-valued Kripke models which extends Definition 4.2.5.

Definition 4.3.1 (Language $\mathcal{L}_{B\Box\alpha}$) *Let* At *be a countable set of atomic propositions. The set* $\mathcal{L}_{B\Box\alpha}$ *of formulas and the set* $\mathcal{AM}_4$ *of four-valued action models are defined by mutual induction by the following grammar:*

$$\mathcal{L}_{B\Box\alpha} \ni \varphi ::= p \mid \mathsf{t} \mid \top \mid \bot \mid \mathsf{f} \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid {\sim}\varphi \mid \varphi \supset \varphi \mid \Box\varphi \mid [\alpha_k]\varphi$$
$$\mathcal{AM}_4 \ni \alpha ::= (\langle K, R_\alpha \rangle, (\varphi_1, ...\varphi_n))$$

where $p \in \mathsf{At}$ *and* $\alpha \in \mathcal{AM}_4$*,* k *is a state in* α*,* $\langle K, R_\alpha\rangle$ *is a finite four-valued Kripke frame, and* $\varphi_1, ..., \varphi_n \in \mathcal{L}_{B\Box\alpha}$*, where* $n = |K|$*.* ⊣

Derived connectives $\otimes, \oplus, \neg, \Diamond\ \rightarrow, *, \leftrightarrow$ are defined as before. Moreover, we let $\langle\alpha_k\rangle\varphi := {\sim}[\alpha_k]{\sim}\varphi$. A pair $\langle\alpha, k\rangle$, written as α_k, is called a *(four-valued) epistemic action*, where $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$ and $k \in K$.

Definition 4.3.2 (Bimodal action model) *A bimodal action model (over* $\mathcal{L}_{B\Box\alpha}$*) is a tuple* $\alpha = \langle K, R^+_\alpha, R^-_\alpha, Pre_\alpha\rangle$ *where* $\langle K, R^+_\alpha, R^-_\alpha\rangle$ *is a bimodal Kripke frame and* $Pre_\alpha : K \to \mathcal{L}_{B\Box\alpha}$ *is the precondition map that assigns a formula* $\varphi \in \mathcal{L}_{B\Box\alpha}$ *to each* $k \in K$*. We denote by* $\mathcal{AM}_2$ *the collection of bimodal action models (over* $\mathcal{L}_{B\Box\alpha}$*).* ⊣

Definition 4.3.3 *Let* $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$ *be a four-valued action model. The bimodal action model associated with* α *is a tuple* $\alpha_{\bowtie} = \langle K, R^+_\alpha, R^-_\alpha, Pre_\alpha\rangle$*, where* $\langle K, R^+_\alpha, R^-_\alpha\rangle = \langle K, R_\alpha\rangle_{\bowtie}$ *(Def. 4.2.21) is the bimodal Kripke frame associated with* $\langle K, R_\alpha\rangle$*.* ⊣

Definition 4.3.4 *Let* $\alpha = \langle K, R^+_\alpha, R^-_\alpha, Pre_\alpha\rangle$ *be a bimodal action model, where* $Pre_\alpha : K \to \mathcal{L}_{B\Box\alpha}$*. The four-valued action model associated with* α *is a tuple* $\alpha^{\bowtie} = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$*, where* $\langle K, R_\alpha\rangle = \langle K, R^+_\alpha, R^-_\alpha\rangle^{\bowtie}$ *(Def. 4.2.22) is the four-valued Kripke frame associated with* α*.*⊣

It follows from (4.6) and (4.7) that for all four-valued action models $\alpha \in \mathcal{AM}_4$, and all bimodal action models $\alpha' \in \mathcal{AM}_2$, $(\alpha_{\bowtie})^{\bowtie} = \alpha$ and $((\alpha')^{\bowtie})_{\bowtie} = \alpha'$.

4.3.1 Algebraic semantics for BAML

Recall from Section 3.4 that epistemic updates on modal Boolean algebras are defined in two steps:

1. We construct the intermediate algebra (Def. 3.4.2).
2. We define the updated algebra as the (pseudo-)quotient of the intermediate algebra (Def. 3.4.3).

Here, we follow an analogous construction, but now on modal bilattices. To this end, we use the correspondence between modal bilattices and bimodal Boolean algebras.

Epistemic updates on bimodal Boolean algebras and bimodal frames

In this part, we apply the methods of [105] to carry out the above steps (1) and (2) for bimodal Boolean algebras and *bimodal action models* over them. We first recall the relevant definitions.

Definition 4.3.5 (Algebraic bimodal action models) *Let $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1}\rangle$ be a bimodal Boolean algebra. A bimodal action model over $\mathbb{A}$ is a tuple $a = \langle K, R_a^+, R_a^-, Pre_a\rangle$ where $\langle K, R_a^+, R_a^-\rangle$ is a bimodal Kripke frame and $Pre_a\colon K \to \mathbb{A}$ is a function that assigns to each state of the action model an element in $\mathbb{A}$.* ⊣

Definition 4.3.6 (Intermediate and updated bimodal frame (model)) *For every bimodal Kripke frame $\mathcal{F} = \langle S, R^+, R^-\rangle$ and every bimodal action model $a = \langle K, R_a^+, R_a^-, Pre_a\rangle$ over the complex algebra of $\mathcal{F}$, the* intermediate bimodal Kripke frame *is defined as*

$$\coprod_a \mathcal{F} = \langle \coprod_K S, R^+ \times R_a^+, R^- \times R_a^-\rangle$$

where $\coprod_K S$ is the $|K|$-fold coproduct of S (which is isomorphic, as a set, to the Cartesian product $S \times K$), and for all $s, s' \in S$, and all $i, j \in K$:

$$(s', j) \in (R^+ \times R_a^+)((s,i)) \quad \textit{iff} \quad s' \in R^+(s) \textit{ and } j \in R_a^+(i)$$

$$(s', j) \in (R^- \times R_a^-)((s,i)) \quad \textit{iff} \quad s' \in R^-(s) \textit{ and } j \in R_a^-(i).$$

The updated bimodal Kripke frame $\mathcal{F}_a$ is defined as the subframe of $\coprod_a \mathcal{F}$ the domain of which is the subset

$$S_a = \{(s,j) \in \coprod_K S \ : \ s \in Pre_a(j)\}.$$

The intermediate bimodal Kripke model *is defined as*

$$\coprod_a \mathcal{M} = \langle \coprod_a \mathcal{F}, \coprod_a V^+, \coprod_a V^-\rangle$$

where $\coprod_a V^+, \coprod_a V^- : \mathsf{At} \to \mathcal{P}(\coprod_K S)$ are valuations and defined as follows: for all $p \in \mathsf{At}$

$$\coprod_a V^+(p)(k) = V^+(p) \quad \textit{and} \quad \coprod_a V^-(p)(k) = V^-(p).$$

The updated bimodal Kripke model *is then defined as the submodel of* $\coprod_a \mathcal{M}$

$$\mathcal{M}_a = \langle \mathcal{F}_a, V_a^+, V_a^- \rangle$$

where $\mathcal{F}_a$ *is the updated bimodal Kripke frame and* $V_a^+, V_a^- : \mathsf{At} \to \mathcal{P}(S_a)$ *are defined as follows: for all* $p \in \mathsf{At}$,

$$V_a^+(p) = \coprod_a V^+(p) \cap S_a \quad \textit{and} \quad V_a^-(p) = \coprod_a V^-(p) \cap S_a.$$

The notion of intermediate algebra for bimodal Boolean algebras is just the bimodal version of the intermediate algebra for modal Boolean algebras (cf. Def. 3.4.2).

Definition 4.3.7 (Intermediate bimodal Boolean algebra) *Let* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1} \rangle$ *be a bimodal Boolean algebra, and let* $a = \langle K, R_a^+, R_a^-, Pre_a \rangle$ *be a bimodal action model over* $\mathbb{A}$. *The intermediate bimodal Boolean algebra*

$$\prod_a \mathbb{A} = \langle A^K, \wedge_{\prod_a \mathbb{A}}, \vee_{\prod_a \mathbb{A}}, \neg_{\prod_a \mathbb{A}} \Box^+_{\prod_a \mathbb{A}}, \Box^-_{\prod_a \mathbb{A}}, \mathbf{0}_{\prod_a \mathbb{A}}, \mathbf{1}_{\prod_a \mathbb{A}} \rangle$$

is the bimodal Boolean algebra, where the Boolean operations are defined pointwise and for every $f \colon K \to \mathbb{A}$, *the functions* $\Box^+_{\prod_a \mathbb{A}} f \colon K \to \mathbb{A}$ *and* $\Box^-_{\prod_a \mathbb{A}} f \colon K \to \mathbb{A}$ *are defined as follows:*

$$(\Box^+_{\prod_a \mathbb{A}} f)(k) = \bigwedge_{k' \in R_a^+(k)} \Box^+_{\mathbb{A}} f(k') \quad \textit{and} \quad (\Box^-_{\prod_a \mathbb{A}} f)(k) = \bigwedge_{k' \in R_a^-(k)} \Box^-_{\mathbb{A}} f(k') \tag{4.12}$$

It is straightforward to check that the intermediate bimodal algebra of a perfect bimodal Boolean algebra is perfect.

As for step (2), we define the updated bimodal algebra via the pseudo-quotient.

Definition 4.3.8 (Updated bimodal algebra) *Let* $\mathbb{A} = \langle A, \wedge, \vee, \neg, \Box^+, \Box^-, \mathbf{0}, \mathbf{1} \rangle$ *be a bimodal Boolean algebra and let* $a = \langle K, R_a^+, R_a^-, Pre_a \rangle$ *be a bimodal action model over* $\mathbb{A}$. *The updated bimodal algebra is the pseudo-quotient of the intermediate bimodal algebra, i.e., it is the bimodal Boolean algebra*

$$\mathbb{A}_a = \langle A^K / \equiv_a, \wedge_a, \vee_a, \neg_a, \Box_a^+, \Box_a^-, \mathbf{0}_a, \mathbf{1}_a \rangle$$

where $\equiv_a$ *is the equivalence relation on* A^K *defined as for modal Boolean algebras (see page 50), that is, for all* $f, g \in A^K$,

$$f \equiv_a g \quad \textit{iff} \quad f \wedge Pre_a = g \wedge Pre_a. \tag{4.13}$$

The modal operators $\Box_a^+$ *and* $\Box_a^-$ *are defined by*

$$\Box_a^+[f] = [\Box^+_{\prod_a \mathbb{A}}(f \to Pre_a)] \quad \textit{and} \quad \Box_a^-[f] = [\Box^-_{\prod_a \mathbb{A}}(f \to Pre_a)]. \tag{4.14}$$

The updated bimodal algebra of a perfect bimodal Boolean algebra is perfect [139].

Proposition 4.3.9 *Let* $\mathcal{F} = \langle S, R^+, R^- \rangle$ *be a bimodal Kripke frame and let* $a = \langle K, R_a^+, R_a^-, Pre_a \rangle$ *be a bimodal action model over the complex algebra* $\mathcal{F}^\bullet$ *of* $\mathcal{F}$. *Then, we have*

$$(\coprod_a \mathcal{F})^\bullet \cong \prod_a \mathcal{F}^\bullet \quad \textit{and} \quad (\mathcal{F}_a)^\bullet \cong (\mathcal{F}^\bullet)_a.$$

Proof The proof is similar to the proof of [105, Prop. 3-1]. □

Proposition 4.3.10 *For every perfect bimodal Boolean algebra* $\mathbb{A}$ *and every bimodal action model* a *over* $\mathbb{A}$,

$$(\prod_a \mathbb{A})_\bullet \cong \coprod_a \mathbb{A}_\bullet \quad \textit{and} \quad (\mathbb{A}_a)_\bullet \cong (\mathbb{A}_\bullet)_a. \tag{4.15}$$

Proof The proof is similar to the proof of [105, Fact 23]. □

Epistemic updates on modal bilattices and four-valued Kripke frames

Similar to (bi)modal Boolean algebras, epistemic updates on modal bilattices is characterised in two steps: (1) we define the intermediate bilattice, then (2) the updated bilattice is obtained as quotient of the intermediate bilattice. We do these two steps using the transformation between bimodal Boolean algebras and modal bilattices along with the definitions of intermediate bimodal algebra and updated bimodal algebra. The following diagram illustrates the construction.

$$
\begin{array}{ccccc}
 & & \text{intermediate bilattice} & & \text{updated bilattice} \\
\mathbb{B} & \twoheadleftarrow & \prod_a \mathbb{B} & \twoheadrightarrow & \mathbb{B}_a \\
(.)^{\bowtie} \uparrow \; \downarrow (.)_{\bowtie} & & (.)^{\bowtie} \uparrow \; \downarrow (.)_{\bowtie} & & (.)^{\bowtie} \uparrow \; \downarrow (.)_{\bowtie} \\
\mathbb{B}_{\bowtie} & \twoheadleftarrow & \prod_{a_{\bowtie}} \mathbb{B}_{\bowtie} & \twoheadrightarrow & (\mathbb{B}_{\bowtie})_{a_{\bowtie}}
\end{array}
$$

We begin by defining action models over modal bilattices.

Definition 4.3.11 (Four-valued action model over a modal bilattice) *Let $\mathbb{B}$ be a modal bilattice. A* four-valued action model *over $\mathbb{B}$ is a tuple $a = \langle K, R_a, Pre_a \rangle$ where $\langle K, R_a \rangle$ is a four-valued Kripke frame, and $Pre_a \colon K \to \mathbb{B}$ in a function that assigns to each state of K an element from $\mathbb{B}$.* ⊣

As we have seen (see Def. 4.2.21), for every four-valued Kripke frame $\mathcal{F}$ there is a bimodal Kripke frame $\mathcal{F}_{\bowtie}$. Also, for every modal bilattice $\mathbb{B}$ there is a bimodal Boolean algebra $\mathbb{B}_{\bowtie}$ (see Equation (4.2)). Hence, for every four-valued action model $a = \langle K, R_a, Pre_a \rangle$ over a modal bilattice $\mathbb{B}$, there is a bimodal action model

$$a_{\bowtie} = \langle K, R^+_{a_{\bowtie}}, R^-_{a_{\bowtie}}, Pre_{a_{\bowtie}} \rangle \tag{4.16}$$

over $\mathbb{B}_{\bowtie}$, where $\langle K, R^+_{a_{\bowtie}}, R^-_{a_{\bowtie}} \rangle$ is the bimodal Kripke frame associated with a (Def. 4.2.21), and $Pre_{a_{\bowtie}} \colon K \to \mathbb{B}_{\bowtie}$ maps each $k \in K$ to $[Pre_a(k)]$ where $[Pre_a(k)]$ is the equivalence class of $Pre_a(k)$ under the relation defined in (4.1). Using these constructions we define the intermediate and updated four-valued Kripke frame.

Definition 4.3.12 (Intermediate and updated four-valued Kripke frame) *For every four-valued Kripke frame $\mathcal{F}$ and every four-valued action model a over the complex bilattice of $\mathcal{F}$, its* intermediate four-valued Kripke frame *is defined as*

$$\coprod_a \mathcal{F} = (\coprod_{a_{\bowtie}} \mathcal{F}_{\bowtie})^{\bowtie}$$

and the updated four-valued Kripke frame *is defined as*

$$\mathcal{F}_a = ((\mathcal{F}_{\bowtie})_{a_{\bowtie}})^{\bowtie}.$$

The intermediate four-valued Kripke model *is defined as*

$$\coprod_a \mathcal{M} = (\coprod_{a_{\bowtie}} \mathcal{M}_{\bowtie})^{\bowtie}$$

Finally, the updated four-valued Kripke model *is defined as*

$$\mathcal{M}_a = ((\mathcal{M}_{\bowtie})_{a_{\bowtie}})^{\bowtie}.$$

The following notion of intermediate bilattices is based on the ideas used in [105].

Definition 4.3.13 (Intermediate modal bilattice) *For every modal bilattice $\mathbb{B} = \langle \mathbf{B}, \Box \rangle$ and all four-valued action models $a = \langle K, R_a, Pre_a \rangle$ over $\mathbb{B}$, the* intermediate modal bilattice *is defined as*

$$\prod_a \mathbb{B} = \langle \mathbf{B}^K, \Box_{\prod_a \mathbb{B}} \rangle$$

where $\mathbf{B}^K$ is the $|K|$-fold direct product of the implicative bilattice $\mathbf{B}$ and the modal operator is given, for each $f \in B^K$ and $j \in K$, by

$$\Box_{\prod_a \mathbb{B}} f(j) = \bigwedge \{ \Box_{\mathbb{B}} f(i) \ : \ i \in K \ \ and \ \ R_a(j,i) \in \{\mathsf{t}, \top\} \}.$$

Using that $\Diamond = {\sim}\Box{\sim}$, it follows that:

$$\Diamond_{\prod_a \mathbb{B}} f(j) = \bigvee \{ \Diamond_{\mathbb{B}} f(i) \ : \ i \in K \ \ and \ \ R_a(j,i) \in \{\mathsf{t}, \top\} \}$$

The next lemma shows that the intermediate modal bilattice of every modal bilattice is a modal bilattice, and shows how one can construct up to isomorphism the intermediate modal bilattice from the intermediate bimodal algebra via the operations $(.)_{\bowtie}$ and $(.)^{\bowtie}$.

Lemma 4.3.14 *Let $a = \langle K, R_a, Pre_a \rangle$ be an action model over a modal bilattice $\mathbb{B}$. Then*

$$\prod_a \mathbb{B} \cong (\prod_{a_{\bowtie}} \mathbb{B}_{\bowtie})^{\bowtie}.$$

Proof Let $\mathbb{B} = \langle \mathbf{B}, \Box \rangle$ be a modal bilattice and $a = \langle K, R_a, Pre_a \rangle$ be a four-valued action model over $\mathbb{B}$. We define the map $h : \prod_a \mathbb{B} \to (\prod_{a_{\bowtie}} \mathbb{B}_{\bowtie})^{\bowtie}$ as follows: for every $f : K \to B$,

$$h(f) = \langle f_1, f_2 \rangle$$

where $f_1, f_2 : K \to B/\equiv$ are such that $f_1(k) = [f(k)]$ and $f_2(k) = [{\sim}f(k)]$ for every $k \in K$ where $[f(k)]$ and $[{\sim}f(k)]$ are the equivalence classes of $f(k)$ and ${\sim}f(k)$ under $\equiv$, the equivalence relation defined in (4.1). It is clear that h is well-defined. Now we check that h is an injective and surjective modal bilattice homomorphism. We first check that h is an injective map. Let $f, g : K \to B$ be such that $h(f) = h(g)$. Then it follows by the definition of h that $\langle f_1, f_2 \rangle = \langle g_1, g_2 \rangle$, which means $f_1(k) = g_1(k)$ and $f_2(k) = g_2(k)$, for all $k \in K$. This implies that $[f(k)] = [g(k)]$ and $[{\sim}f(k)] = [{\sim}g(k)]$, i.e., $f(k) \wedge g(k) = f(k) \oplus g(k)$ and ${\sim}f(k) \wedge {\sim}g(k) = {\sim}f(k) \oplus {\sim}g(k)$. Then we have that

$$\begin{aligned}
f(k) \vee g(k) = {\sim}{\sim}(f(k) \vee g(k)) &= {\sim}({\sim}f(k) \wedge {\sim}g(k)) & \\
&= {\sim}({\sim}f(k) \oplus {\sim}g(k)) & ([{\sim}f(k)] = [{\sim}g(k)]) \\
&= {\sim}{\sim}f(k) \oplus {\sim}{\sim}g(k) & ({\sim}(x \oplus y) = {\sim}x \oplus {\sim}y) \\
&= f(k) \oplus g(k) & ({\sim}{\sim}x = x) \\
&= f(k) \wedge g(k). & ([f(k)] = [g(k)])
\end{aligned}$$

Hence, we obtain that $f(k) \wedge g(k) = f(k) \vee g(k)$ for all $k \in K$, which implies that $f(k) = g(k)$ for all $k \in K$. Therefore, h is a injection. Now, we show that h is surjective. Let $f_1, f_2 : K \to B/\equiv$ be such that $f_1(k) = [g_1(k)]$ and $f_2(k) = [g_2(k)]$, for all $k \in K$, where $g_1, g_2 \in B^K$. We will show that there is $g \in B^K$ such that $h(g) = \langle f_1, f_2 \rangle$. We define $g : K \to B$ by $g(k) =$

$(g_1(k) \wedge \top) \otimes (\sim g_2(k) \vee \top)$. Then, by [137, Proposition 2.1.6][1], it follows that $h(g) = \langle f_1, f_2 \rangle$. Hence, h is surjective. Using a similar line of argumentation as in the proof of [140, Theorem 5.18], one can show that h preserves modal bilattice operations. □

The following result shows that the notion of intermediate bilattice is dual to the notion of intermediate four-valued Kripke frame (Definition 4.3.13).

Theorem 4.3.15 *Let $\mathcal{F} = \langle S, R \rangle$ be a four-valued Kripke frame and let $a = \langle K, R_a, Pre_a \rangle$ be a four-valued action model over the complex bilattice $\mathcal{F}^*$. Then the intermediate modal bilattice $\prod_a \mathcal{F}^*$ is isomorphic to the complex bilattice of the intermediate four-valued Kripke frame. In other words:*

$$\prod_a \mathcal{F}^* \cong (\coprod_a \mathcal{F})^*.$$

Proof We first compute the complex bilattice of the intermediate four-valued Kripke frame $\mathcal{F}$, i.e., $(\coprod_a \mathcal{F})^*$.

$$\begin{aligned}
(\coprod_a \mathcal{F})^* &= (((\coprod_a \mathcal{F})_{\bowtie})^{\bullet})^{\bowtie} && \text{(Def. 4.2.31)}\\
&= ((((\coprod_{a_{\bowtie}} \mathcal{F}_{\bowtie})^{\bowtie})_{\bowtie})^{\bullet})^{\bowtie} && \text{(Def. 4.3.12)}\\
&\cong ((\coprod_{a_{\bowtie}} \mathcal{F}_{\bowtie})^{\bullet})^{\bowtie} && (4.7)\\
&\cong (\prod_{a_{\bowtie}} (\mathcal{F}_{\bowtie})^{\bullet})^{\bowtie} && \text{(Prop. 4.3.9)}
\end{aligned}$$

On the other hand, by Def. 4.2.31 and (4.3) we have that $(\mathcal{F}^*)_{\bowtie} \cong (\mathcal{F}_{\bowtie})^{\bullet}$. Hence,

$$\prod_{a_{\bowtie}} (\mathcal{F}^*)_{\bowtie} \cong \prod_{a_{\bowtie}} (\mathcal{F}_{\bowtie})^{\bullet},$$

and thus using Lemma 4.3.14,

$$\prod_a \mathcal{F}^* \cong (\prod_{a_{\bowtie}} (\mathcal{F}^*)_{\bowtie})^{\bowtie} \cong (\prod_{a_{\bowtie}} (\mathcal{F}_{\bowtie})^{\bullet})^{\bowtie} \cong (\coprod_a \mathcal{F})^*.$$

□

At this point we can apply a similar definition of pseudo-quotient[2] from [138, 139] to obtain a suitable notion of quotient of an intermediate bilattice.

Definition 4.3.16 (Pseudo-quotient relation) *Given a four-valued action model $a = \langle K, R_a, Pre_a \rangle$ over a modal bilattice $\mathbb{B} = \langle \mathbf{B}, \Box \rangle$, we define the relation $\equiv_a$ on $\prod_a \mathbb{B}$ as follows: for all $f, g \in \prod_a \mathbb{B}$*

$$f \equiv_a g \quad \textit{iff} \quad f \wedge \neg\neg Pre_a = g \wedge \neg\neg Pre_a. \tag{4.17}$$

[1] We note that Proposition 2.1.6 in [137] is about interlaced bilattices. Since distributive bilattices are interlaced bilattices, this result also holds for our case.

[2] Note that in [138, 139], $\sim$ denotes the Boolean negation which we denote by $\neg$.

Because $Pre_a \in \prod_a \mathbb{B}$ and $\wedge, \neg$ are algebraic operations of $\prod_a \mathbb{B}$, it follows from [138, Fact 2.2] that $\equiv_a$ is a congruence of the non-modal reduct $\mathbf{B}^K$ of $\prod_a \mathbb{B}$. We will denote the equivalence class of $f \in \prod_a \mathbb{B}$ by $[f]_a$ (or simply by $[f]$ when there is no risk of confusion) and the quotient set $\mathbf{B}^K/\equiv_a$ by $\mathbb{B}_a$.

Proposition 4.3.17 *Let $\mathbb{B}$ be a modal bilattice, and let $a = \langle K, R_a, Pre_a \rangle$ be a four-valued action model over $\mathbb{B}$. Then the following hold:*

(i) $[f \wedge \neg\neg Pre_a] = [f]$ *for every* $f \in \prod_a \mathbb{B}$. *Hence, for every* $f \in \prod_a \mathbb{B}$, *there exists a unique* $g \in \prod_a \mathbb{B}$ *such that* $g \in [f]$ *and* $g \leq_t \neg\neg Pre_a$.

(ii) *For all* $f, g \in \prod_a \mathbb{B}$, *we have* $[f] \leq_t [g]$ *iff* $f \wedge \neg\neg Pre_a \leq_t g \wedge \neg\neg Pre_a$. ⊣

Proof *Item (i)* follows from that fact that $\wedge$ is idempotent, hence:

$$(f \wedge \neg\neg Pre_a) \wedge \neg\neg Pre_a = f \wedge \neg\neg Pre_a$$

for all $f \in \prod_a \mathbb{B}$. This proves the first part of the statement and the existence claim of the second part. As for the uniqueness, if there is $g \in \prod_a \mathbb{B}$ such that $g \in [f]$ and $g \leq_t \neg\neg Pre_a$, then we have that $g = g \wedge \neg\neg Pre_a = f \wedge \neg\neg Pre_a$, as since $g \in [f]$.

Item (ii). The right to left direction follows from the definition of $\equiv_a$. For the other direction, first we note that if $[f] \leq_t [g]$ then there are $f', g' \in \prod_a \mathbb{B}$ such that $f' \in [f]$, $g' \in [g]$, and $f' \leq_t g'$. Then we have $f' \wedge \neg\neg Pre_a \leq g' \wedge \neg\neg Pre_a$, as desired.

□

The modalities on the pseudo-quotient can now be introduced as another application of the definitions in [138, 139]. For every action model $a = \langle K, R_a, Pre_a \rangle$ over a modal bilattice $\mathbb{B}$ and every $f \in \prod_a \mathbb{B}$, we let

$$\Box_a[f] := [\Box_{\prod_a \mathbb{B}}(f \to \neg\neg Pre_a)].$$

The dual operator is given by $\Diamond_a[f] := \neg\Box_a\neg[f]$.

Proposition 4.3.18 *Let $\mathbb{B}$ be a modal bilattice, and let $a = \langle K, R_\alpha, Pre_a \rangle$ be an action model over $\mathbb{B}$. Then*

(i) the algebra $\mathbb{B}_a$ is a modal bilattice.

(ii) $\mathbb{B}_a \cong ((\mathbb{B}_{\bowtie})_{a_{\bowtie}})^{\bowtie}$ ⊣

Proof *Item (i)* follows from a similar line of argumentation as in the proof of Fact 2.4. in [138].

Item (ii). Define the isomorphism

$$\begin{aligned} h' \colon \mathbb{B}_a &\to ((\mathbb{B}_{\bowtie})_{a_{\bowtie}})^{\bowtie} \\ [f]_a &\mapsto \langle [f_1]_{a_{\bowtie}}, [f_2]_{a_{\bowtie}} \rangle, \end{aligned}$$

where $[f]_a$ is the equivalence class of f under $\equiv_a$ (Def. 4.3.16) and $[f_1]_{a_{\bowtie}}$ and $[f_2]_{a_{\bowtie}}$ are the equivalence classes of $f_1, f_2 \in \prod_{a_{\bowtie}} \mathbb{B}_{\bowtie}$ under $\equiv_{a_{\bowtie}}$ (Def. 4.3.8) such that $f_1(k) = [f(k)]$ and $f_2(k) = [\sim f(k)]$ in which $[f(k)]$ and $[\sim f(k)]$ are the equivalence classes of $f(k)$ and $\sim f(k)$ under the equivalence relation $\equiv$ defined in (4.1). The proof that shows h' is a modal bilattice isomorphism is similar to the proof of Lemma 4.3.14. □

Definition 4.3.19 (Updated modal bilattice) *Given a modal bilattice $\mathbb{B}$ and a four-valued action model a over $\mathbb{B}$, the modal bilattice $\mathbb{B}_a$ is called the updated modal bilattice.* ⊣

Proposition 4.3.20 *Let $\mathcal{F} = \langle S, R\rangle$ be a four-valued Kripke frame, and let $a = \langle K, R_a, Pre_a\rangle$ be a four-valued action model over the complex bilattice of $\mathcal{F}$. Then,*

$$(\mathcal{F}_a)^* \cong (\mathcal{F}^*)_a \tag{4.18}$$

Proof It follows from the definition of complex bilattices and of updates on four-valued Kripke frames and modal bilattices together with (4.3). □

Item (i) of Proposition 4.3.17 says that each $\equiv_a$-equivalence class has a canonical representative, which is the unique element in the given class which is below the element $\neg\neg Pre_a$ in the truth order. Hence we can define an (injective) map $i' \colon \mathbb{B}_a \to \prod_a \mathbb{B}$ by

$$i'([f]) = f \wedge \neg\neg Pre_a \tag{4.19}$$

for all $[f] \in \mathbb{B}_a$. Denoting by $q \colon \prod_a \mathbb{B} \to \mathbb{B}_a$ the canonical quotient map, we have that the composition $q \circ i'$ is the identity on $\mathbb{B}_a$.

The map $i' \colon \mathbb{B}_a \to \prod_a \mathbb{B}$ plays a key role in the definition of interpretation of **BAML** formulas on algebraic models. In the next theorem we characterise i' in terms of the inclusion map $\mathsf{in} \colon \mathcal{F}_a \to \coprod_a \mathcal{F}$ (see Def. 4.3.12).

Lemma 4.3.21 *Let $\mathcal{F}$ be a four-valued Kripke frame and let $a = \langle K, R_a, Pre_a\rangle$ be a four-valued action model over $\mathcal{F}^*$. Then the following diagram commutes:*

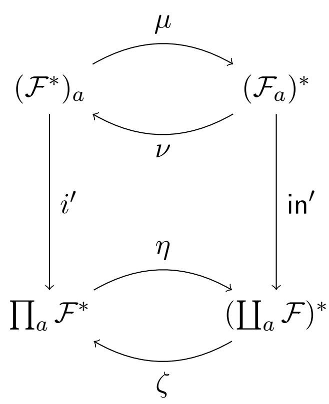

where $\mathsf{in}' = \mathsf{in}[.] \times \mathsf{in}[.]$, $\mu \colon (\mathcal{F}^*)_a \to (\mathcal{F}_a)^*$ and $\nu \colon (\mathcal{F}_a)^* \to (\mathcal{F}^*)_a$ are the modal bilattice isomorphisms from Prop. 4.3.20 such that $\mu^{-1} = \nu$, and $\eta : \prod_a \mathcal{F}^* \to (\coprod_a \mathcal{F})^*$ and $\zeta : (\coprod_a \mathcal{F})^* \to \prod_a \mathcal{F}^*$ are the modal bilattice isomorphisms from Thm. 4.3.15 such that $\eta^{-1} = \zeta$. ⊣

Proof Apply [105, Proposition 3.6] to the bimodal Boolean algebra $((\mathcal{F}^*)_a)_{\bowtie}$ and use the correspondence between bimodal Boolean algebras and modal bilattices. □

Definition 4.3.22 (Bilattice model) *A* bilattice model of **BAML** *is a tuple $\mathcal{B} = \langle \mathbb{B}, V\rangle$ such that $\mathbb{B}$ is a modal bilattice and $V \colon \mathsf{At} \to \mathbb{B}$. For every bilattice model $\mathcal{B}$ and every action model $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$. Let $a = \langle K, R_a, Pre_a\rangle$ be the action model over $\mathbb{B}$ induced by α via the unique extension of V to $\mathcal{L}_{B\Box\alpha}$, $[\![-]\!]_{\mathcal{B}} : \mathcal{L}_{B\Box\alpha} \to \mathbb{B}$ with $Pre_a(k) = [\![Pre_\alpha(k)]\!]_{\mathcal{B}}$, for every $k \in K$. We let $\prod_\alpha \mathcal{B} = \langle \prod_\alpha \mathbb{B}, \prod_\alpha V\rangle$, where $(\prod_\alpha V)(p) := \prod_a V(p)$ for every $p \in \mathsf{At}$. Likewise, we define the updated bilattice model as $\mathcal{B}_a := \langle \mathbb{B}_a, V_a\rangle$ where $V_a = q \circ \prod_\alpha V$ and $q : \prod_a \mathbb{B} \to \mathbb{B}_a$ is the quotient map.* ⊣

Definition 4.3.23 (Algebraic semantics for BAML) *Given a bilattice model $\mathcal{B} = \langle \mathbb{B}, V\rangle$, the extension map $[\![.]\!]_{\mathcal{B}} : \mathcal{L}_{B\Box\alpha} \to \mathbb{B}$ is defined as follows:*

$$
\begin{aligned}
[\![p]\!]_{\mathcal{B}} &:= V(p) \\
[\![c]\!]_{\mathcal{B}} &:= c_{\mathbb{B}} && \text{for } c \in \{\mathsf{f}, \mathsf{t}, \bot, \top\} \\
[\![\sim\varphi]\!]_{\mathcal{B}} &:= \sim_{\mathbb{B}} [\![\varphi]\!]_{\mathcal{B}} \\
[\![\Box\varphi]\!]_{\mathcal{B}} &:= \Box_{\mathbb{B}} [\![\varphi]\!]_{\mathcal{B}} \\
[\![\varphi \bullet \psi]\!]_{\mathcal{B}} &:= [\![\varphi]\!]_{\mathcal{B}} \bullet_{\mathbb{B}} [\![\psi]\!]_{\mathcal{B}} && \text{for } \bullet \in \{\wedge, \vee, \otimes, \oplus \supset\} \\
[\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{B}} &:= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{B}} \wedge_{\mathbb{B}} (\pi_k \circ i')([\![\varphi]\!]_{\mathcal{B}_a}) \\
[\![[\alpha_k]\varphi]\!]_{\mathcal{B}} &:= [\![Pre_\alpha(k)]\!]_{\mathcal{B}} \supset_{\mathbb{B}} (\pi_k \circ i')([\![\varphi]\!]_{\mathcal{B}_a})
\end{aligned}
$$

where $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$ and $k \in K$, $i' : \mathbb{B}_a \to \prod_a \mathbb{B}$ is defined in (4.19) and $\pi_k : \prod_\alpha \mathbb{B} \to \mathbb{B}$ is the projection onto the k-th coordinate. ⊣

For a set Γ of $\mathcal{L}_{B\Box\alpha}$-formulas, we write $\Gamma \models_{\mathbf{BAML}} \varphi$ if for every bilattice model $\mathcal{B} = \langle \mathbb{B}, V\rangle$ and every bifilter $F \subseteq B$, we have that $[\![\gamma]\!]_{\mathcal{B}} \in F$ for all $\gamma \in \Gamma$ implies $[\![\varphi]\!]_{\mathcal{B}} \in F$. A formula φ is valid, if for every bilattice model $\mathcal{B} = \langle \mathbb{B}, V\rangle$, $[\![\varphi]\!]_{\mathcal{B}} \geq \top_{\mathbb{B}}$.

4.3.2 Relational semantics for BAML

We will now use this algebraic semantics and duality theory to introduce a relational semantics for **BAML** based on four-valued Kripke models. The epistemic updates on four-valued Kripke models is obtained via duality from epistemic updates on bimodal Kripke models and correspondence between bimodal Kripke models and four-valued Kripke models.

Epistemic updates on bimodal Kripke models

The following constructions are analogous to that of Def. 4.3.6. The only difference is that here we use bimodal action models in which the preconditions are defined over $\mathcal{L}_{B\Box\alpha}$ whereas in Def. 4.3.6 we use bimodal action models in which the preconditions are defined over the complex algebras of bimodal Kripke frames.

Definition 4.3.24 (α-intermediate and α-update bimodal Kripke frame (model)) *Let $\mathcal{F} = \langle S, R^+, R^-\rangle$ be a bimodal Kripke frame and let $\mathcal{M} = \langle \mathcal{F}, V^+, V^-\rangle$ be a bimodal Kripke model based on $\mathcal{F}$, and let $\alpha = \langle K, R^+, R^-, Pre_\alpha\rangle$ be a bimodal action model. The* α-intermediate bimodal Kripke frame *is defined in the same way as in Def. 4.3.6, i.e.,*

$$\coprod_\alpha \mathcal{F} := \coprod_a \mathcal{F} \tag{4.20}$$

where $a = \langle K, R^+, R^-, Pre_a\rangle$ is a bimodal action model over the complex algebra of $\mathcal{F}$ induced by α via the unique extension of V^+. The α-updated bimodal Kripke frame *is defined as*

$$\mathcal{F}_\alpha = \mathcal{F}_a \tag{4.21}$$

The α-intermediate bimodal Kripke model *is then defined as the coproduct structure*

$$\coprod_\alpha \mathcal{M} = \coprod_a \mathcal{M} \tag{4.22}$$

Finally, we define the α-updated bimodal Kripke model

$$\mathcal{M}_\alpha := \mathcal{M}_a. \tag{4.23}$$

We can now define a relational semantics of **BAML** with respect to bimodal Kripke models.

Definition 4.3.25 (2Mdl-semantics of BAML) *Let* $\mathcal{M} = \langle S, R^+, R^-, V^+, V^- \rangle$ *be a bimodal Kripke model,* $s \in S$*. Then*

$$\begin{array}{lll}
(\mathcal{M}, s) \models c & \textit{iff} & c \in \{\mathsf{t}, \top\} \\
(\mathcal{M}, s) \models p & \textit{iff} & s \in V^+(p) \\
(\mathcal{M}, s) \models \varphi \wedge \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ and } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \varphi \vee \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ or } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models {\sim}\varphi & \textit{iff} & (\mathcal{M}, s) \not\models \varphi \\
(\mathcal{M}, s) \models \varphi \supset \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ implies } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \Box\varphi & \textit{iff} & \textit{for all } t \in S : t \in R^+(s) \textit{ then } (\mathcal{M}, t) \models \varphi \\
(\mathcal{M}, s) \models [\alpha_k]\varphi & \textit{iff} & (\mathcal{M}, s) \models Pre_\alpha(k) \textit{ implies } (\mathcal{M}_{\alpha_{\bowtie}}, (s, k)) \models \varphi.
\end{array}$$

where $\alpha = \langle K, R_\alpha, Pre_\alpha \rangle \in \mathcal{AM}_4$ *is a four-valued action model, and* $\alpha_{\bowtie} = \langle K, R^+_{\alpha_{\bowtie}}, R^-_{\alpha_{\bowtie}}, Pre_\alpha \rangle \in \mathcal{AM}_2$ *is the bimodal action model associated with* α *(Def. 4.3.3). We denote the truth set of* φ *in* $\mathcal{M}$ *by* $[\![\varphi]\!]_{\mathcal{M}} := \{s \in S : (\mathcal{M}, s) \models \varphi\}$. ⊣

Epistemic updates on four-valued Kripke models

The following constructions are analogous to that of Def. 4.3.12. The only difference is that here we use four-valued action models in which the preconditions are defined over $\mathcal{L}_{B\Box\alpha}$ whereas in Def. 4.3.12 we use four-valued action models in which the preconditions are defined over the complex bilattice of four-valued Kripke frames.

Definition 4.3.26 (α-intermediate and α-updated four-valued Kripke frame (model)) *Let* $\mathcal{F} = \langle S, R \rangle$ *be a four-valued Kripke frame, and let* $\alpha = \langle K, R, Pre_\alpha \rangle \in \mathcal{AM}_4$*. The* α-intermediate four-valued Kripke frame *is*

$$\coprod_\alpha \mathcal{F} = (\coprod_{\alpha_{\bowtie}} \mathcal{F}_{\bowtie})^{\bowtie}, \tag{4.24}$$

and the α-updated four-valued Kripke frame *is*

$$\mathcal{F}_\alpha = ((\mathcal{F}_{\bowtie})_{\alpha_{\bowtie}})^{\bowtie}, \tag{4.25}$$

where $\alpha_{\bowtie} = \langle K, R^+_\alpha, R^-_\alpha, Pre_\alpha \rangle$ *is the bimodal action model associated with* α*. We note that by Def. 4.3.24,*

$$\coprod_\alpha \mathcal{F} = \coprod_a \mathcal{F},$$

where $\coprod_a \mathcal{F}$ *is defined in 4.3.12. Then, the* α-intermediate four-valued Kripke model *is defined as*

$$\coprod_\alpha \mathcal{M} = (\coprod_{\alpha_{\bowtie}} \mathcal{M}_{\bowtie})^{\bowtie} \tag{4.26}$$

and define the α-updated four-valued Kripke model *as*

$$\mathcal{M}_\alpha = ((\mathcal{M}_{\bowtie})_{\alpha_{\bowtie}})^{\bowtie}, \tag{4.27}$$

where $\mathcal{M}_{\bowtie}$ *and* $\alpha_{\bowtie}$ *are the bimodal Kripke model, and bimodal action model associated with* $\mathcal{M}$ *and* α *(Def. 4.2.21).* ⊣

We can now define a relational semantics for **BAML** with respect to four-valued Kripke models.

Definition 4.3.27 (fourMdl-semantics of BAML) *Let $\mathcal{M} = \langle S, R, V\rangle$ be a four-valued Kripke model and $s \in S$. Then*

$$\begin{array}{lll}
(\mathcal{M}, s) \models c & \textit{iff} & c \in \{\mathsf{t}, \top\} \\
(\mathcal{M}, s) \models p & \textit{iff} & V(p, s) \in \{\mathsf{t}, \top\} \\
(\mathcal{M}, s) \models \varphi \wedge \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ and } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \varphi \vee \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ or } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \sim\varphi & \textit{iff} & (\mathcal{M}, s) \not\models \varphi \\
(\mathcal{M}, s) \models \varphi \supset \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ implies } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \Box\varphi & \textit{iff} & \textit{for all } t \in S : R(s,t) \in \{\mathsf{t}, \top\} \textit{ implies } (\mathcal{M}, t) \models \varphi \\
(\mathcal{M}, s) \models [\alpha_k]\varphi & \textit{iff} & (\mathcal{M}, s) \models Pre_\alpha(k) \textit{ implies } (\mathcal{M}_\alpha, (s,k)) \models \varphi
\end{array}$$

where $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$, and $k \in K$. ⊣

The next proposition shows that the two relational semantics of **BAML** with respect to four-valued Kripke models, and with respect to bimodal Kripke models are equivalent.

Proposition 4.3.28 *Let $\varphi \in \mathcal{L}_{B\Box\alpha}$ be an arbitrary formula. Then,*

1. *For every four-valued Kripke model $\mathcal{M} = \langle S, R, V\rangle$, and state $s \in S$,*

$$(\mathcal{M}, s) \models \varphi \quad \textit{iff} \quad (\mathcal{M}_{\bowtie}, s) \models \varphi.$$

where $\mathcal{M}_{\bowtie}$ is a bimodal Kripke model defined in Def. 4.2.21.

2. *For every bimodal Kripke model $\mathcal{M} = \langle S, R^+, R^-, V^+, V^-\rangle$, and state $s \in S$,*

$$(\mathcal{M}, s) \models \varphi \quad \textit{iff} \quad (\mathcal{M}^{\bowtie}, s) \models \varphi.$$

where $\mathcal{M}^{\bowtie}$ is a four-valued Kripke model defined in Def. 4.2.22. ⊣

Proof We only prove item 1, and leave item 2 to the reader. We first define the following order on formulas and action models: Let $\varphi, \psi \in \mathcal{L}_{B\Box\alpha}$ and $\alpha = \langle K, R_\alpha, Pre_\alpha\rangle \in \mathcal{AM}_4$. Then

$$\begin{array}{ll}
\psi < \varphi & \text{if } \psi \text{ is a subformula of } \varphi \\
\varphi < \alpha & \text{if } \varphi = Pre_\alpha(k), \text{ for some } k \in K \\
\alpha < \varphi & \text{if } \varphi = [\alpha_k]\psi, \text{ for some } \psi \in \mathcal{L}_{B\Box\alpha}, \text{ and for some } k \in K.
\end{array}$$

We denote by $\bar{<}$, the transitive closure of $<$.

The proof is then by well-founded induction on the relation $\bar{<}$. We must show that for all $\varphi \in \mathcal{L}_{B\Box\alpha}$ and all four-valued Kripke models $\mathcal{M} = \langle S, R, V\rangle$ and $s \in S$,

$$(\text{For all } \psi \bar{<} \varphi : (\mathcal{M}, s) \models \psi \text{ iff } (\mathcal{M}_{\bowtie}, s) \models \psi) \implies ((\mathcal{M}, s) \models \varphi \text{ iff } (\mathcal{M}_{\bowtie}, s) \models \varphi)$$

So, suppose $\varphi \in \mathcal{L}_{B\Box\alpha}$ and for all $\psi \in \mathcal{L}_{B\Box\alpha}$ with $\psi \bar{<} \varphi$, and all four-valued Kripke models $\mathcal{M} = \langle S, R, V\rangle$, $(\mathcal{M}, s) \models \psi$ iff $(\mathcal{M}_{\bowtie}, s) \models \psi$. We show that $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}_{\bowtie}, s) \models \varphi$. We distinguish between the different forms that φ can have. Cases for logical constants, the propositional variables and the bilattice connectives are elementary. Now consider the case where $\varphi := \Box\psi$. Let $\mathcal{M}$ be a four-valued Kripke model and s in a state in $\mathcal{M}$. Then we have

$$\begin{aligned}
(\mathcal{M}, s) \models \Box\psi \quad &\iff \quad \text{for all } t \in S : R(s,t) \in \{\mathsf{t}, \top\} \text{ implies } (\mathcal{M}, t) \models \psi \\
&\iff \quad \text{for all } t \in S : t \in R^+(s) \text{ implies } (\mathcal{M}, t) \models \psi \qquad \text{(Def. 4.2.21, Eq. 4.4)} \\
&\iff \quad \text{for all } t \in S : t \in R^+(s) \text{ implies } (\mathcal{M}_{\bowtie}, t) \models \psi \qquad \text{(Ind. hyp. , } \psi \bar{<} \varphi) \\
&\iff \quad (\mathcal{M}_{\bowtie}, s) \models \Box\psi.
\end{aligned}$$

Now assume that $\varphi := [\alpha_k]\psi$. We first note that since $\psi \bar{<} \varphi$, by the induction hypothesis,

$$(\mathcal{M}_\alpha, (s,j)) \models \psi \quad \iff \quad ((\mathcal{M}_\alpha)_{\bowtie}, (s,j)) \models \psi. \tag{4.28}$$

where $(s,j) \in S_\times$. By the definition of $\mathcal{M}_\alpha$ in (4.27) and (4.6), it follows that $(\mathcal{M}_\alpha)_{\bowtie} \cong (\mathcal{M}_{\bowtie})_{\alpha_{\bowtie}}$, hence by the fact that isomorphic models satisfy the same formulas, and (4.28), we obtain that

$$(\mathcal{M}_\alpha, (s,j)) \models \psi \quad \iff \quad ((\mathcal{M}_{\bowtie})_{a_{\bowtie}}, (s,j)) \models \psi. \tag{4.29}$$

Then we have

$$\begin{aligned}
(\mathcal{M}, s) \models [\alpha_k]\psi \quad &\iff \quad (\mathcal{M}, s) \models Pre_\alpha(k) \text{ implies } (\mathcal{M}_\alpha, (s,k)) \models \psi \\
\text{(Ind. hyp. and } Pre_\alpha(k) \bar{<} \varphi) \quad &\iff \quad (\mathcal{M}_{\bowtie}, s) \models Pre_\alpha(k) \text{ implies } ((\mathcal{M}_{\bowtie})_{a_{\bowtie}}, (s,j)) \models \psi \\
&\iff \quad (\mathcal{M}_{\bowtie}, s) \models [\alpha_k]\psi.
\end{aligned}$$

□

The next proposition shows that the mechanism of epistemic updates remains essentially unchanged when moving from the classical to a bilattice setting.

Proposition 4.3.29 *For every perfect modal bilattice $\mathbb{B}$ and every four-valued action model $a = \langle K, R_a, Pre_a \rangle$ over $\mathbb{B}$, we have*

1. $(\prod_a \mathbb{B})_* \cong \coprod_a \mathbb{B}_*$, *and*
2. $(\mathbb{B}_a)_* \cong (\mathbb{B}_*)_a$. ⊣

Proof *Item (1).* Let $\mathbb{B}$ be a modal bilattice and a be four-valued action model over $\mathbb{B}$. We have

$$\begin{aligned}
(\prod_a \mathbb{B})_* &\cong ((\prod_{a_{\bowtie}} \mathbb{B}_{\bowtie})^{\bowtie})_* && \text{(Lemma 4.3.14)} \\
&\cong (((\prod_{a_{\bowtie}} \mathbb{B}_{\bowtie})_\bullet)^{\bowtie} && \text{(Def. 4.2.31, (4.3))}
\end{aligned}$$

On the other hand,

$$\begin{aligned}
\coprod_a \mathbb{B}_* &= (\coprod_{a_{\bowtie}} (\mathbb{B}_*)_{\bowtie})^{\bowtie} && \text{(4.24)} \\
&\cong (\coprod_{a_{\bowtie}} (\mathbb{B}_{\bowtie})_\bullet)^{\bowtie} && \text{(Def. 4.2.31, (4.7))}
\end{aligned}$$

Hence, by Prop. 4.3.10, it follows that $(\prod_a \mathbb{B})_* \cong \coprod_a \mathbb{B}_*$.
Item (2).

$$\begin{aligned}(\mathbb{B}_a)_* &\cong (((\mathbb{B}_{\bowtie})_{a_{\bowtie}})^{\bowtie})_* && \text{(Prop. 4.3.18(2))}\\ &\cong (((\mathbb{B}_{\bowtie})_{a_{\bowtie}})_{\bullet})^{\bowtie} && \text{(Def. 4.2.31, (4.3))}\end{aligned}$$

On the other hand,

$$\begin{aligned}(\mathbb{B}_*)_a &= (((\mathbb{B}_*)_{\bowtie})_{a_{\bowtie}})^{\bowtie} && \text{(Def. 4.3.12)}\\ &\cong (((\mathbb{B}_{\bowtie})_{\bullet})_{a_{\bowtie}})^{\bowtie} && \text{(Def. 4.2.31, (4.7))}\\ &\cong (((\mathbb{B}_{\bowtie})_{a_{\bowtie}})_{\bullet})^{\bowtie} && \text{(4.3.10)}\end{aligned}$$

Hence, $(\mathbb{B}_a)_* \cong (\mathbb{B}_*)_a$. □

4.4 Axiomatisation

In this section we give a Hilbert-style proof system for **BAML** on the class of four-valued frames. We show that it is sound and complete. The proof system $\mathbb{BAML}$ for the logic **BAML** consists of all the rules and axioms given in the Tables 4.3, 4.4, and 4.5. Table 4.3 contains the propositional part, Table 4.4 contains the (static) modal part, and Table 4.5 contains the dynamic (modal) part. The rules and principles of Table 4.5 for **BAML** resemble those for bilattice public announcement logic introduced in [138, 139], the only difference is the rule **RE**. In the rule **RE**, called 'replacement of equivalents', $\chi[\varphi/p]$ means uniform substitution of all occurrences of p in χ by φ (this can be easily defined inductively). The derivation rule 'replacement of equivalents' (**RE**) was erroneously missing in previous axiomatizations of non-classical dynamic epistemic logics [114, 105, 138, 139]. In the absence of (**RE**), the reduction strategy of $\mathbb{BAML}$ to its static fragment, as sketched in the proof of Theorem 4.4.7, later, would not succeed.[3]

The axioms in Table 4.5 are the expected reduction rules for any logical structure following a dynamic modality for action model execution. Clearly, as in **BAML** we have constants, we have axioms for the reduction of each of those constants. But there is nothing surprising about them. The other axioms may look more familiar to the reader informed about dynamic epistemic logics, except for the occasional need of the $\neg\neg$ binding of preconditions Pre_α: this is to ensure the restriction of the possible values of $\neg\neg Pre_\alpha$, namely to $\mathtt{t}$ and $\mathtt{f}$ only.

The axiom $\langle\alpha_k\rangle p \leftrightarrow (\neg\neg Pre_\alpha(k) \wedge p)$, called ($\langle\boldsymbol{\alpha_k}\rangle$**-atoms**), guarantees that the value of atoms is preserved after update. Such an axiom is often formulated both for positive and for negative atoms (i.e., for literals). The axiom for negative atoms is indeed a theorem of our axiomatization. We show its derivation as an example.

Example 4.4.1 *$\langle\alpha_k\rangle{\sim}p \leftrightarrow \neg\neg Pre_\alpha(k) \wedge {\sim}p$ is a theorem of $\mathbb{BAML}$.*

$$\begin{aligned}&(1)\langle\alpha_k\rangle{\sim}p \leftrightarrow \neg\neg Pre_\alpha(k) \wedge {\sim}\langle\alpha_k\rangle p && (\langle\boldsymbol{\alpha_k}\rangle{\sim})\\ &(2)\langle\alpha_k\rangle{\sim}p \leftrightarrow \neg\neg Pre_\alpha(k) \wedge {\sim}(\neg\neg Pre_\alpha(k) \wedge p) && (\text{(1)}, (\langle\boldsymbol{\alpha_k}\rangle\textbf{-atoms}), \textbf{RE})\\ &(3)\langle\alpha_k\rangle{\sim}p \leftrightarrow \neg\neg Pre_\alpha(k) \wedge ({\sim}(\neg\neg Pre_\alpha(k)) \vee {\sim}p) && (\text{(2)}, ({\sim}\wedge), \textbf{RE})\\ &(4){\sim}(\neg\neg Pre_\alpha(k)) \leftrightarrow \neg Pre_\alpha(k) && (({\sim}\supset))\end{aligned}$$

[3]The rule **RE** is needed because we use an inside-out reduction strategy. For the alternative outside-in reduction strategy, **RE** is not needed, but then one needs a reduction axiom of shape "$\langle\alpha_k\rangle\langle\beta\rangle\varphi \leftrightarrow \ldots$" as well as a rule "from $\varphi \to \psi$ infer $\langle\alpha_k\rangle\varphi \to \langle\alpha_k\rangle\psi$" ($\alpha$-monotonicity). For classical dynamic epistemic logics, for the special case of public announcement logics, such variants are discussed in detail in [157].

(**$\langle\alpha_k\rangle$-constants**)	$\langle\alpha_k\rangle\mathsf{f} \leftrightarrow \mathsf{f}$ $\quad \langle\alpha_k\rangle\mathsf{t} \leftrightarrow \neg\neg Pre_\alpha(k)$
	$\langle\alpha_k\rangle\top \leftrightarrow (Pre_\alpha(k) \wedge \top)$ $\quad \langle\alpha_k\rangle\bot \leftrightarrow \sim(Pre_\alpha(k) \supset \bot)$
(**$\langle\alpha_k\rangle$-atoms**)	$\langle\alpha_k\rangle p \leftrightarrow (\neg\neg Pre_\alpha(k) \wedge p)$
(**$\langle\alpha_k\rangle\wedge$**)	$\langle\alpha_k\rangle(\varphi \wedge \psi) \leftrightarrow (\langle\alpha_k\rangle\varphi \wedge \langle\alpha_k\rangle\psi)$
(**$\langle\alpha_k\rangle\vee$**)	$\langle\alpha_k\rangle(\varphi \vee \psi) \leftrightarrow (\langle\alpha_k\rangle\varphi \vee \langle\alpha_k\rangle\psi)$
(**$\langle\alpha_k\rangle\supset$**)	$\langle\alpha_k\rangle(\varphi \supset \psi) \leftrightarrow (\neg\neg Pre_\alpha(k) \wedge (\langle\alpha_k\rangle\varphi \supset \langle\alpha_k\rangle\psi))$
(**$\langle\alpha_k\rangle\sim$**)	$\langle\alpha_k\rangle\sim\varphi \leftrightarrow (\neg\neg Pre_\alpha(k) \wedge \sim\langle\alpha_k\rangle\varphi))$
(**$\langle\alpha_k\rangle\Diamond$**)	$\langle\alpha_k\rangle\Diamond\varphi \leftrightarrow (\neg\neg Pre_\alpha(k) \wedge \bigvee\{\Diamond\langle\alpha_j\rangle\varphi \ : \ R_\alpha(k,j) \in \{\mathsf{t}, \top\}\})$
(**RE**)	from $\varphi \leftrightarrow \psi$ infer $\chi[\varphi/p] \leftrightarrow \chi[\psi/p]$

Table 4.5. The axiomatization $\mathbb{BAML}$ for the logic **BAML** consists of all rules and axioms of $\mathbb{BML}$ (see Tables 4.3 and 4.4) and the above axioms and rule.

(5)$\langle\alpha_k\rangle\sim p \leftrightarrow \neg\neg Pre_\alpha(k) \wedge (\neg Pre_\alpha(k) \vee \sim p)$ ((3), (4), **RE**)

(6)$\langle\alpha_k\rangle\sim p \leftrightarrow (\neg\neg Pre_\alpha(k) \wedge \neg Pre_\alpha(k)) \vee (\neg\neg Pre_\alpha(k) \wedge \sim p)$ ((5), distributivity)

(7)$(\neg\neg Pre_\alpha(k) \wedge \neg Pre_\alpha(k)) \leftrightarrow \mathsf{f}$ (**LB**)

(8)$\langle\alpha_k\rangle\sim p \leftrightarrow \mathsf{f} \vee (\neg\neg Pre_\alpha(k) \wedge \sim p)$ ((6), (7), **RE**)

(9)$\langle\alpha_k\rangle\sim p \leftrightarrow \neg\neg Pre_\alpha(k) \wedge \sim p$ ((8), $f \vee \varphi \leftrightarrow \varphi$)

We now proceed by showing soundness and completeness. The following lemmas are needed to establish that $\mathbb{BAML}$ is sound with respect to the algebraic semantics. Most proofs are straightforward adaptations of the lemmas from [138, 139].

Lemma 4.4.2 ([138], Lemma 4.1) *Let $\mathcal{M} = \langle\mathbb{B}, V\rangle$ be a bilattice model and φ a formula such that $[\![\varphi]\!]_{\mathcal{M}_\alpha} = q([\![\varphi]\!]_{\prod_\alpha \mathcal{M}})$ for any four-valued action α. Then:*

- $[\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge [\![\varphi]\!]_{\mathcal{M}}$
- $[\![[\alpha_k]\varphi]\!]_{\mathcal{M}} = [\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset [\![\varphi]\!]_{\mathcal{M}}$ ⊣

Lemma 4.4.3 ([138], Fact 4.2) *Let $\mathbb{B}$ be a modal bilattice and a be a four-valued action model over $\mathbb{B}$, and let $i' : \mathbb{B}_a \to \prod_a \mathbb{B}$ be given by $[g] \mapsto g \wedge \neg\neg Pre_a$. Then for every $[b], [c] \in \mathbb{B}_a$:*

- $i'([b] \wedge [c]) = i'([b]) \wedge i'([c])$;
- $i'([b] \vee [c]) = i'([b]) \vee i'([c])$;
- $i'([b] \supset [c]) = \neg\neg Pre_a \wedge (i'([b]) \supset i'([c]))$;
- $i'(\sim[b]) = \neg\neg Pre_a \wedge \sim i'([b])$;
- $i'(\Box_a[b]) = \neg\neg Pre_a \wedge \Box_{\prod_a \mathbb{B}}(Pre_a \supset i'([b]))$;
- $i'(\Diamond_a[b]) = \neg\neg Pre_a \wedge \Diamond_{\prod_a \mathbb{B}}(i'([b]) \wedge \neg\neg Pre_a)$. ⊣

Lemma 4.4.4 ([138, Lemma 4.3]) *Let $\mathcal{M} = \langle\mathbb{B}, V\rangle$ be a bilattice model with underlying modal bilattice $\mathbf{B} = \langle B, \wedge, \vee, \supset, \sim, \Diamond, \Box, \mathsf{f}, \mathsf{t}, \bot, \top\rangle$. For every four-valued action model α and all formulas φ and ψ:*

1. $[\![\langle\alpha_k\rangle(\varphi \vee \psi)]\!]_{\mathcal{M}} = [\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}} \vee [\![\langle\alpha_k\rangle\psi]\!]_{\mathcal{M}}$;
2. $[\![\langle\alpha_k\rangle(\varphi \wedge \psi)]\!]_{\mathcal{M}} = [\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}} \wedge [\![\langle\alpha_k\rangle\psi]\!]_{\mathcal{M}}$;
3. $[\![\langle\alpha_k\rangle(\varphi \supset \psi)]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge ([\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}} \supset [\![\langle\alpha_k\rangle\psi]\!])$;
4. $[\![\langle\alpha_k\rangle{\sim}\varphi]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge {\sim}[\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}}$;
5. $[\![[\alpha_k]\varphi]\!]_{\mathcal{M}} = [\![{\sim}\langle\alpha_k\rangle{\sim}\varphi]\!]_{\mathcal{M}}$;
6. $[\![\langle\alpha_k\rangle\Diamond\varphi]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee\{\Diamond_{\mathbf{B}}([\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\}$;
7. $[\![\langle\alpha_k\rangle\Box\varphi]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigwedge\{\Box_{\mathbb{B}}([\![[\alpha_j]\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\}$. $\dashv$

Proof Items (1)-(5) can be prove using the same line of arguments as in the proof of [138, Lemma 4.3]. Here, we only show items (6) and (7). Concerning (6), first observe that:

$$\begin{aligned}
&\pi_k \circ i'([\![\Diamond\varphi]\!]_{\mathcal{M}_\alpha}) = \pi_k(\neg\neg Pre_\alpha \wedge \Diamond_{\prod_\alpha \mathbf{B}}(\neg\neg Pre_\alpha \wedge i'([\![\varphi]\!]_{\mathcal{M}_\alpha})))\\
&= \neg\neg Pre_a(k) \wedge \bigvee \left\{\Diamond_{\mathbb{B}}(\neg\neg Pre_a(j) \wedge i'([\![\varphi]\!]_{\mathcal{M}_\alpha}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee \left\{\Diamond_{\mathbb{B}}(\neg\neg Pre_a(j)) \wedge i'([\![\varphi]\!]_{\mathcal{M}_\alpha}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee \left\{\Diamond_{\mathbb{B}}(\neg\neg[\![Pre_\alpha(j)]\!]_{\mathcal{M}} \wedge \pi_j \circ i'([\![\varphi]\!]_{\mathcal{M}_\alpha}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee \left\{\Diamond_{\mathbb{B}}(\neg\neg[\![Pre_\alpha(j)]\!]_{\mathcal{M}} \wedge \pi_j \circ i'([\![\varphi]\!]_{\mathcal{M}\alpha_j}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee \left\{\Diamond_{\mathbb{B}}([\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}.
\end{aligned}$$

To justify the equality between lines 4 and 5 above, note that $\mathcal{M}_\alpha$ is independent from the point of α, i.e., $(\mathcal{M}_\alpha =)\ \mathcal{M}_{\alpha_k} = \mathcal{M}_{\alpha_j}$. Then:

$$\begin{aligned}
&[\![\langle\alpha_k\rangle\Diamond\varphi]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \pi_k \circ i'([\![\Diamond\varphi]\!]_{\mathcal{M}_\alpha}\varphi)\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee \left\{\Diamond_{\mathbf{B}}([\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee \left\{\Diamond_{\mathbf{B}}([\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}
\end{aligned}$$

To show item (7), we preliminarily observe that

$$\begin{aligned}
&\pi_k \circ i'([\![\Box\varphi]\!]_{\mathcal{M}_\alpha}) = \pi_k(\neg\neg Pre_a(k) \wedge \Box_{\prod_\alpha \mathbb{B}}(Pre_a \supset i'([\![\varphi]\!]_{\mathcal{M}_\alpha})))\\
&= \neg\neg Pre_a(k) \wedge \bigwedge \left\{\Box_{\mathbb{B}}(Pre_a(j) \supset i'([\![\varphi]\!]_{\mathcal{M}_\alpha}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset \bigwedge \left\{\Box_{\mathbb{B}}(Pre_a(j) \supset i'([\![\varphi]\!]_{\mathcal{M}_\alpha}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigwedge\left\{\Box_{\mathbb{B}}([\![Pre_\alpha(j)]\!]_{\mathcal{M}} \supset \pi_j \circ i'([\![\varphi]\!]_{\mathcal{M}_\alpha}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigwedge\left\{\Box_{\mathbb{B}}([\![Pre_\alpha(j)]\!]_{\mathcal{M}} \supset \pi_j \circ i'([\![\varphi]\!]_{\mathcal{M}\alpha_j}))(j) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigwedge\left\{\Box_{\mathbb{B}}([\![[\alpha_j]\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}.
\end{aligned}$$

Hence:

$$\begin{aligned}
&[\![\langle\alpha_k\rangle\Box\varphi]\!]_{\mathcal{M}} = \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \pi_k \circ i'([\![\Box\varphi]\!]_{\mathcal{M}_\alpha})\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge (\neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigwedge \left\{\Box_{\mathbb{B}}([\![[\alpha_j]\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\})\\
&= \neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigwedge \left\{\Box_{\mathbb{B}}([\![[\alpha_j]\varphi]\!]_{\mathcal{M}}) \,:\, R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}.
\end{aligned}$$

$\Box$

Item (v) of Lemma 4.4.4 justifies our usage of $[\alpha_k]\varphi$ as an abbreviation for $\sim\langle\alpha_k\rangle\sim\varphi$.

The next lemma is also helpful for the intuition linking the relational and algebraic setting, but is not strictly necessary in the completeness proof, wherein we use that $\langle\alpha_k\rangle\varphi$ is a primitive language construct and $[\alpha_k]\varphi$ a derived one.

Lemma 4.4.5 ([138, Fact 4.4]) *Let $\mathcal{M} = \langle \mathbb{B}, V\rangle$ be a bilattice model with underlying modal bilattice $\mathbb{B} = \langle B, \wedge, \vee, \supset, \sim, \Box, \mathsf{f}, \mathsf{t}, \bot, \top\rangle$. For every action model α and all formulas φ and ψ in $\mathcal{L}_{B\Box\alpha}$:*

1. $[\![[\alpha_k](\varphi \wedge \psi)]\!]_{\mathcal{M}} = [\![[\alpha_k]\varphi]\!]_{\mathcal{M}} \wedge [\![[\alpha_k]\psi]\!]_{\mathcal{M}}$
2. $[\![[\alpha_k](\varphi \vee \psi)]\!]_{\mathcal{M}} = [\![[\alpha_k]Pre_\alpha]\!]_{\mathcal{M}} \supset ([\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}} \vee [\![\langle\alpha_k\rangle\psi]\!]_{\mathcal{M}})$
3. $[\![[\alpha_k](\varphi \supset \psi)]\!]_{\mathcal{M}} = [\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}} \supset [\![\langle\alpha_k\rangle\psi]\!]_{\mathcal{M}}$
4. $[\![[\alpha_k]\sim\varphi]\!]_{\mathcal{M}} = \sim[\![\langle\alpha_k\rangle\varphi]\!]_{\mathcal{M}}$
5. $[\![[\alpha_k]\varphi]\!]_{\mathcal{M}} = [\![\sim\langle\alpha_k\rangle\sim\varphi]\!]_{\mathcal{M}}$
6. $[\![[\alpha_k]\Diamond\varphi]\!]_{\mathcal{M}} = [\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset \bigvee\left\{\Diamond_{\mathbb{B}}([\![[\alpha_j]\varphi]\!]_{\mathcal{M}}) \;:\; R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}$
7. $[\![[\alpha_k]\Box\varphi]\!]_{\mathcal{M}} = [\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset \bigwedge\left\{\Box_{\mathbb{B}}([\![[\alpha_j]\varphi]\!]_{\mathcal{M}}) \;:\; R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}$ ⊣

Proof Item (7) follows from Lemma 4.4.4.vi. The only other item of interest is (6):

$$\begin{aligned}
[\![[\alpha_k]\Diamond\varphi]\!]_{\mathcal{M}} &= [\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset i'([\![\Diamond\varphi]\!]_{\mathcal{M}_\alpha}) \\
&= [\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset (\neg\neg[\![Pre_\alpha(k)]\!]_{\mathcal{M}} \wedge \bigvee\left\{\Diamond_{\mathbb{B}}([\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{M}}) \;:\; R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\}) \\
&= [\![Pre_\alpha(k)]\!]_{\mathcal{M}} \supset \bigvee\left\{\Diamond_{\mathbb{B}}([\![\langle\alpha_j\rangle\varphi]\!]_{\mathcal{M}}) \;:\; R_\alpha(k,j) \in \{\mathsf{t}, \top\}\right\} \qquad (*)
\end{aligned}$$

Equivalence $(*)$ holds since in every modal bilattice we have that $x \supset \neg\neg x = \mathsf{t}$ and that $(x \supset y) \wedge (x \supset z) = x \supset (y \wedge z)$. □

Lemma 4.4.6 *The rule* **RE** *is sound: if $\models \varphi \leftrightarrow \psi$ then $\models \chi[\varphi/p] \leftrightarrow \chi[\psi/p]$.* ⊣

Proof Let $\varphi, \psi \in \mathcal{L}_{BL}$ be such that $\vDash \varphi \leftrightarrow \psi$. We will prove that for all $\chi \in \mathcal{L}_{BL}$, and any Kripke model $\mathcal{M} = \langle S, R, V\rangle$ and state $s \in S$:

$$[\![\chi[\varphi/p], s]\!]_{\mathcal{M}} = [\![(\chi[\psi/p], s)]\!]_{\mathcal{M}}.$$

where $[\![-,-]\!] \colon \mathcal{L}_{B\Box\alpha} \times S \to \mathsf{FOUR}$ is the unique extension of the valuation V to $\mathcal{L}_{B\Box\alpha}$. Since every four-valued Kripke model is isomorphic to a bimodal Kripke model, and the value of formulas is preserved under isomorphism [25, Prop. 2.9(ii)], we prove the theorem for bimodal Kripke models. So, let $\mathcal{M} = \langle S, R^+, R^-, V^+, V^-\rangle$ be a bimodal Kripke model. We prove by induction on the structure of χ that

$$[\![\chi[\varphi/p], s]\!]^+_{\mathcal{M}} = [\![(\chi[\psi/p], s)]\!]^+_{\mathcal{M}} \quad \text{and} \quad [\![\chi[\varphi/p], s]\!]^-_{\mathcal{M}} = [\![(\chi[\psi/p], s)]\!]^-_{\mathcal{M}} \tag{4.30}$$

where $[\![.]\!]^+ : \mathcal{L}_{B\Box\alpha} \to \mathcal{P}(S)$ and $[\![.]\!]^- : \mathcal{L}_{B\Box\alpha} \to \mathcal{P}(S)$ are the extensions of V^+ and V^- to $\mathcal{L}_{B\Box\alpha}$, respectively.

— The case where χ is a logical constant or an atomic proposition is immediate.

— If $\chi = \gamma \bullet \delta$, where $\bullet \in \{\wedge, \vee, \supset\}$, or $\chi = \sim\gamma$, use that V is a homomorphism in its first argument with respect to bilattice operators.
— If $\chi = \Box\gamma$, then

$$\begin{aligned} s \in [\![(\Box\gamma[\varphi/p]]\!]^+ &\iff \text{for all } s' \in S : s' \in R^+(s) \text{ implies } s \in [\![\gamma[\varphi/p]]\!]^+ \\ &\iff \text{for all } s' \in S : s' \in R^+(s) \text{ implies } s \in [\![\gamma[\psi/p]]\!]^+ \qquad \text{(Indictive hyp.)} \\ &\iff s \in [\![(\Box\gamma[\psi/p]]\!]^+. \end{aligned}$$

A similar argument shows that $[\![\Box\gamma[\varphi/p]]\!]^- = [\![\Box\gamma[\psi/p]]\!]^-$.

Finally, let $\chi = \langle\alpha_k\rangle\gamma$. We show that $[\![(\langle\alpha_k\rangle\gamma[\varphi/p])]\!]^+ = [\![(\langle\alpha_k\rangle\gamma[\psi/p])]\!]^+$ and that $[\![\langle\alpha_k\rangle\gamma[\varphi/p]]\!]^- = [\![\langle\alpha_k\rangle\gamma[\psi/p]]\!]^-$. Let $\mathcal{M}_{\alpha_{\bowtie}} = \langle S_\times, R^+_\times, R^-_\times, V^+_\times, V^-_\times\rangle$. By inductive hypothesis, for every $(s,k) \in S_\times$, $[\![\gamma[\varphi/p]]\!]^+_{\mathcal{M}_{\alpha_{\bowtie}}} = [\![\gamma[\psi/p]]\!]^+_{\mathcal{M}_{\alpha_{\bowtie}}}$, $[\![\gamma[\varphi/p]]\!]^-_{\mathcal{M}_{\alpha_{\bowtie}}} = [\![\gamma[\psi/p]]\!]^-_{\mathcal{M}_{\alpha_{\bowtie}}}$, and $Pre_\alpha(k)[\varphi/p] = Pre_\alpha(k)[\psi/p]$, where $[\![.]\!]^+_{\mathcal{M}_{\alpha_{\bowtie}}}, [\![.]\!]^-_{\mathcal{M}_{\alpha_{\bowtie}}} : \mathcal{L}_{B\Box\alpha} \to \mathcal{P}(S_\times)$ are the extensions of $V^+_\times$ and $V^-_\times$, respectively. Then, for every $s \in S$ we have

$$\begin{aligned} s \in [\![\langle\alpha_k\rangle\gamma[\varphi/p]]\!]^+_{\mathcal{M}} &\iff s \in [\![Pre_\alpha(k)[\varphi/p]]\!]^+_{\mathcal{M}} \text{ and } s \in \iota_k^{-1}(\mathsf{in}([\![\gamma[\varphi/p]]\!]^+_{\mathcal{M}_{\alpha_{\bowtie}}})) \\ \text{(Inductive hyp.)} &\iff s \in [\![Pre_\alpha(k)[\psi/p]]\!]^+_{\mathcal{M}} \text{ and } s \in \iota_k^{-1}(\mathsf{in}([\![\gamma[\psi/p]]\!]^+_{\mathcal{M}_{\alpha_{\bowtie}}})) \\ &\iff s \in [\![\langle\alpha_k\rangle\gamma[\psi/p]]\!]^+_{\mathcal{M}}. \end{aligned}$$

A similar argument shows that $[\![\gamma[\varphi/p]]\!]^-_{\mathcal{M}_{\alpha_{\bowtie}}} = [\![\gamma[\psi/p]]\!]^-_{\mathcal{M}_{\alpha_{\bowtie}}}$.

□

We now get to the announced completeness result.

Theorem 4.4.7 *The proof system* $\mathbb{BAML}$ *is sound and complete with respect to algebraic and relational models.* ⊣

Proof The soundness of the preservation of logical constants and propositional variables follows from Lemma 4.4.3. The soundness of the remaining axioms is proved in Lemma 4.4.4. The soundness of **RE** is proved in Lemma 4.4.6.

The proof of completeness is analogous to that of classical and intuitionistic **AML**, and follows from the reducibility of **BAML** to bilattice modal logic.

Let φ be valid. Let us assume that we only use primitive connectives of $\mathcal{L}$ (so, for example, $\langle\alpha_k\rangle$ but not $[\alpha_k]$). Consider some innermost occurrence $\langle\alpha_k\rangle\psi$ of a dynamic modality in φ, where ψ is in the static language. The axioms of $\mathbb{BAML}$ make it possible to transform $\langle\alpha_k\rangle\psi$ into an equivalent formula without a dynamic modality:

We 'push' the dynamic modality down the generation tree of the formula, through the static connectives, until it binds a proposition letter or a constant symbol. There, the dynamic modality disappears, thanks to an application of the appropriate axiom preserving proposition letters or constants, and, crucially, applying the **RE** rule (we *replace* a subformula in a larger expression by an equivalent formula without the dynamic modality).

This process is repeated for all the dynamic modalities of φ, so as to obtain a formula φ' which is provably equivalent to φ. Since φ is valid by assumption, and since provable equivalence preserves validity, by soundness we can conclude that φ' is valid. By Theorem 4.2.13, we can conclude that φ' is a theorem in bilattice modal logic and thus in $\mathbb{BAML}$. Therefore, as φ and φ' are provably equivalent, φ is also a theorem. This concludes the proof.

□

4.5 Case study: Knowledge of inconsistency and incompleteness

A good candidate for a recipient of possibly inconsistent information is the database. You are Hendrik Edeling, a breeder of tulips. Consider the database D1-Acuminata containing information on the colour of a particular tulip that is a candidate for selective breeding. It may contain the information that the tulip is red, or that it is not red, or it may lack this information, or it may, inconsistently so, contain the information that it is both red and not red. In other words, the proposition p *for 'the tulip is red' can have one of the four values* $\mathsf{t}, \mathsf{f}, \top, \bot$. *Let us now consider the perspective of Edeling wishing to consult the database. And let us assume that Edeling is uncertain which of the four states* $\mathsf{t}, \mathsf{f}, \top, \bot$ *the database is in, with respect to the proposition* p. *That makes four possible worlds that he is unable to distinguish. If he now queries the database and get 'yes' as an answer to the query '*p*?', he can rule out two of these four possibilities and keep the worlds wherein* p *has the value* t *and the value* $\top$. *So this is a way to process a* public announcement of the proposition p. *Now a further query to narrow down the options would be querying the database on the value of* $\sim p$, *or, more properly said, querying it on the falsity of* p. *A confirmation that* p *is false reduces Hendrik Edeling's uncertainty because the only remaining world satisfying it, is the one where the value of* p *is* $\top$. *In another sense, Hendrik has become more uncertain again, because he has confirmation that the database is inconsistent. We could also have communicated directly (in one formula) to Edeling that the database is inconsistent. Or that it is consistent, or that it is incomplete (value* $\bot$*). How? Please read on.*

Given initial uncertainty about p, *Edeling may also have to interact with his colleague Sara Burgerhart, another renowned tulip expert. Maybe even a competitor! Consider the action of Burgerhart being informed that the database is lacking information on* p *(the database is incomplete), while Edeling remains uncertain whether she gets this information.*

The information that the agents receive may also be modal. Suppose that Hendrik is being told that $p \wedge \sim\Box p$*: "The tulip is red but you don't know this!" Unlike in two-valued modal logic, this formula may remain true after its announcement. It need not be an unsuccessful update. How come? Again, please read on.*

Bilattice modal logic We model the D1-Acuminata database containing information on that tulip as a world. The proposition that the tulip is red is p. There are four possible worlds. We use mnemonic names for the worlds: $\mathsf{p}_\bot$ is the world where $V(p) = \bot$, p_t is the world where $V(p) = \mathsf{t}$, p_f is the world where $V(p) = \mathsf{f}$, and $\mathsf{p}_\top$ is the world where $V(p) = \top$. Uncertainty about the four worlds is represented by the following model $\mathcal{M}$. The box enclosing the worlds means that they are indistinguishable (the accessibility relation R is the universal relation on this domain) for Hendrik Edeling.

$$\mathcal{M}: \boxed{\mathsf{p}_\bot \quad \mathsf{p}_\mathsf{f} \quad \mathsf{p}_\mathsf{t} \quad \mathsf{p}_\top}$$

We can now evaluate, for example, that $(\mathcal{M}, \mathsf{p}_\mathsf{t}) \models p$, or that $(\mathcal{M}, \mathsf{p}_\top) \models p$ (we recall that $(\mathcal{M}, s) \models \varphi$ means that $V(\varphi, s) \in \{\mathsf{t}, \top\}$). We do not have that $(\mathcal{M}, \mathsf{p}_\mathsf{t}) \models \Box p$, as both p and $\sim p$ are considered possible. Hendrik is uncertain about p. A public announcement $p!$ restricts the model to the p_t and $\mathsf{p}_\top$ state.

$$\mathcal{M}: \boxed{\mathsf{p}_\bot \quad \mathsf{p}_\mathsf{f} \quad \mathsf{p}_\mathsf{t} \quad \mathsf{p}_\top} \quad \overset{p!}{\Rightarrow} \quad \mathcal{M}_p: \boxed{\mathsf{p}_\mathsf{t} \quad \mathsf{p}_\top}$$

A public announcement is a singleton action model with reflexive access. Instead of writing that α is a public announcement of φ we write $\varphi!$; and for the corresponding model update we write $\mathcal{M}_\varphi$ instead of $\mathcal{M}_\alpha$. We can justify the restriction to t and $\top$ by considering this semantics of announcement to be the response to a query $p?$. In both cases the answer will be 'yes'. In two-valued public announcement logic, we are used to having $[p!]\Box p$ as a validity for propositional variables. This is no longer the case in our setting. In particular, for every state s in model $\mathcal{M}$, $(\mathcal{M}, s) \not\models [p!]\Box p$. We recall the semantics of $\Box$ (Def. 4.2.4), namely

$$[\![\Box\varphi, s]\!]_{\mathcal{M}} := \bigwedge\{R(s, s') \to [\![\varphi, s']\!]_{\mathcal{M}} \ : \ s' \in S\},$$

where $\bigwedge$ denotes the infinitary version of $\wedge$ and $\to$ is the strong implication. For the sake of a smooth presentation, we identify the valuation V and its unique extension $[\![-,-]\!]_{\mathcal{M}}$. So far, our models have two-valued accessibility relations, i.e., $(s, s') \in R$ or $(s, s') \notin R$ for all pairs in $\mathcal{M}$. This means that $\Box p$ takes the value $\bigwedge\{\mathsf{t} \to V(p, s) : s \in \{\mathsf{p}_\bot, \mathsf{p}_\mathsf{t}, \mathsf{p}_\mathsf{f}, \mathsf{p}_\top\}\}$. As $V(p, \mathsf{p}_\top) = \top$, and $\mathsf{t} \to \top = \mathsf{f}$ (the other three values are t), $\Box p$ is therefore *false* (in all states of $\mathcal{M}$) and not true. The intuition behind this is that in bilattice modal logic $\Box\varphi$ *is false if* φ *is false in an accessible world.* It is not necessarily the case that $\Box\varphi$ *is true if* φ *is true in all accessible worlds.* In fact, if in one or more of those accessible worlds φ has the value $\top$ (as in our example model $\mathcal{M}$), then φ is also false in an accessible world, and thus we are done for. From $(\mathcal{M}_p, s) \not\models \Box p$ then follows, using that $((\mathcal{M}, s) \models [\alpha_k]\varphi$ iff $((\mathcal{M}, s) \models \mathit{Pre}_\alpha(k)$ implies $(\mathcal{M}_\alpha, (s, k)) \models \varphi))$, that $(\mathcal{M}, s) \not\models [p!]\Box p$.

]Now consider the announcement of $p \wedge {\sim}\Box p$. This formula is known as the Moore sentence [123, 47]. In two-valued public announcement logic, as the result of truthfully announcing it, it becomes false; $[(p \wedge {\sim}\Box p)!]{\sim}(p \wedge {\sim}\Box p)$ is valid in public announcement logic. It is not valid in **BAML**. Similarly to above, we have:

$\mathcal{M}$: $\boxed{\mathsf{p}_\bot \quad \mathsf{p}_\mathsf{f} \quad \mathsf{p}_\mathsf{t} \quad \mathsf{p}_\top}$ $\overset{(p\wedge{\sim}\Box p)!}{\Rightarrow}$ $\mathcal{M}_{p\wedge{\sim}\Box p}$: $\boxed{\mathsf{p}_\mathsf{t} \quad \mathsf{p}_\top}$

Thus, because in $\mathcal{M}_{p\wedge{\sim}\Box p}$ we have that $R(\mathsf{p}_\mathsf{t}, \mathsf{p}_\top) = \mathsf{t}$ and that $V(p, \mathsf{p}_\top) = \top$, it follows that $\mathcal{M}_{p\wedge{\sim}\Box p}, \mathsf{p}_\mathsf{t} \not\models \Box p$. In fact, we now have that $\mathcal{M}_{p\wedge{\sim}\Box p}, \mathsf{p}_\mathsf{t} \models p \wedge {\sim}\Box p$ and thus the (from a modal logical perspective) surprising result that:

$[(p \wedge {\sim}\Box p)!](p \wedge {\sim}\Box p)$ *is satisfiable in* **BAML**.

Having seen some simple examples of announcements and of formulas, and modal formulas, let us present some simple announcements on the status quo of a database, with regard to p.

- the database is consistent: announcement of $\neg(p \wedge {\sim}p)$
- the database is inconsistent: announcement of $p \wedge {\sim}p$
- the database is complete: announcement of $p \vee {\sim}p$
- the database is incomplete: announcement of $\neg(p \vee {\sim}p)$

The four-valued truth tables of these formulas are illustrative.

$\neg$	$(p$	$\wedge$	$\sim$	$p)$
t	$\bot$	$\bot$	$\bot$	$\bot$
t	f	f	t	f
t	t	f	f	t
f	$\top$	$\top$	$\top$	$\top$
$\Uparrow$				

p	$\wedge$	$\sim$	p
$\bot$	$\bot$	$\bot$	$\bot$
f	f	t	f
t	f	f	t
$\top$	$\top$	$\top$	$\top$
	$\Uparrow$		

p	$\vee$	$\sim$	p
$\bot$	$\bot$	$\bot$	$\bot$
f	t	t	f
t	t	f	t
$\top$	$\top$	$\top$	$\top$
	$\Uparrow$		

$\neg$	$(p$	$\vee$	$\sim$	$p)$
t	$\bot$	$\bot$	$\bot$	$\bot$
f	f	t	t	f
f	t	t	f	t
f	$\top$	$\top$	$\top$	$\top$
$\Uparrow$				

Again, we do not necessarily have that after these announcements, the formulas of the announcement are known: $[\sim(p \vee \neg p)!]\Box\sim(p \vee \neg p)$ and $[\neg(p \wedge \sim p)!]\Box\sim(p \vee \neg p)$ are *valid*, but $[(p \vee \sim p)!]\Box(p \vee \sim p)$ and $[(p \wedge \sim p)!]\Box(p \wedge \sim p)$ are *invalid*. (Although $[\neg\neg(p \vee \sim p)!]\Box\neg\neg(p \vee \sim p)$ and $[\neg\neg(p \wedge \sim p)!]\Box\neg\neg(p \vee \sim p)$ are valid.)

It is illustrative to see announcements as answers of queries to the database. When Hendrik queries the database with p? then the answer he gets will be 'yes' if the state of the database is t or $\top$, when he queries the database with $\sim p$? then the answer he gets will be 'yes' if the state of the database is f or $\top$. This is like Fitting's Rosencrantz and Guildenstern (R and G) setting in [66]. In question-answer analysis in two-valued logic [81], a question induces a *partition* on the domain, and a yes/no question, such as a question φ? on the truth of φ, a *dichotomy*. Fitting's Rosencrantz and Guildenstern other answer is 'no'. That is, for either of them, a classical dichotomy. However, it is tempting to see a *question in four-valued logic* differently, namely as inducing (a set of subsets that is) a *partial cover* of the domain. It is a *cover*, as two subsets may have non-empty intersection (namely when they contain worlds where φ has the value $\top$). It is *partial*, as some worlds may not be in any subset, namely when φ has the value $\bot$. If the world has no information on φ (value $\bot$), then 'there is no answer' or, differently said, the answer is: "I don't know." This becomes like the introductory example where you were trying to find your way to the railway station in Nancy. That example also serves to illustrate another, we think, interesting feature of four-valued question-answer analysis: if the value of φ is $\top$, then the answer to the question ?φ is 'yes' (so not 'yes and no'); whereas the answer to the question $\neg\varphi$? is also 'yes'. Knowledge of inconsistency is a higher order feature for a database: whereas consulting memory directly is more straightforward: if you already have the answer 'yes', why trying to rule out the answer 'no'? In other words, questions become *leading questions*. We do not know if this analysis of questions in four-valued logics is common in inquisitive semantics [81].

Roles in dynamic epistemics *To understand dynamic epistemics, also on bilattices, it is important to distinguish different roles: (i) the agent/object/process identified with a propositional variable (the holder of the information), (ii) the agent being uncertain about the proposition, and (iii) the provider of reliable new information (on the proposition), the dynamic part. In our tulip example we have distinguished (i) (the database) from (ii) (Hendrik Edeling), but not (i) from (iii) (the database is queried and provides the answers). In the railway station example (i) and (iii) are separate: accidental pedestrians perform the role of (iii). It is common to view the source of new information, the 'announcing agent', as an anonymous oracle or trusted authority (Hendrik Edeling's system manager, so to speak). In the tulip example we can even think of the different roles as different components of 'the database' as hardware: (i) is RAM, (ii) is the CPU, and (iii) is the interface. In multi-agent examples (where each agent* a *has her own* $\Box^i$ *in the logical language) it is also easier to separate roles.*

Truth values or possible worlds? *If we see Hendrik Edeling as the database, we can consider the value of p his uncertainty. Initially the value of p is* $\bot$. *It then* changes *into* t *once Hendrik gathers positive information on p, and may further change into* $\top$ *if he additionally receives negative information on p. These are so-called* factual (ontic) changes. *But if we see Hendrik as* *different from the database, then his uncertainty is between four worlds of a Kripke model, where a world represents the fixed value of p in the database. Receiving information now means restricting this model in order to finally find out the true value of p. This is* informational (epistemic) change. *The former is quite different from the latter. Factual change can also be modelled in dynamic epistemics, but is outside our scope.*

Multi-agent knowledge and actions Our framework generalizes to a multi-agent modal setting, wherein instead of the modality $\Box$ we have modalities $\Box^i$, for each agent i. The accessibility relation that interprets such a knowledge modality is required to be an equivalence relation (which, so far, is two-valued; four-valued accessibility relations will be considered next). Other scenarios are conceivable, for example for belief, intentions or obligations, or time (with temporal modalities).

Hendrik Edeling has a colleague Sara Burgerhart who is another expert on Acuminata tulips and who may also have access to the same database. We model some scenarios and give typical formulas. Elementary checks on their adequacy are left to the reader. The accessibility relation (equivalence classes) for Hendrik are solid boxes and for Sara they are dashed boxes. Modality $\Box^h$ represents Hendrik's knowledge and $\Box^s$ represents Sara's knowledge.

- Sara knows that the tulip is red.

$\mathsf{p_t} \models \Box^s p$

- Sara knows whether the tulip is red. Hendrik is uncertain whether she knows that. (And we should now add; "and they are both aware of this scenario." We will refrain from doing so from now on.) Sara says to Hendrik: "I know that the tulip is red."

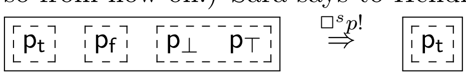

$\mathsf{p_t} \models \sim\Box^h(\Box^s p \vee \Box^s{\sim}p) \wedge [\Box^s p!]\Box^h p$

- Sara knows whether the tulip is red. Hendrik is uncertain whether she knows that. Sara says to Hendrik: "I know whether the tulip is red."

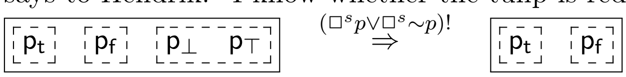

($\mathsf{p_f}$ or) $\mathsf{p_t} \models \sim\Box^h(\Box^s p \vee \Box^s{\sim}p) \wedge [(\Box^s p \vee \Box^s{\sim}p)!](\sim\Box^h p \wedge \Box^h(\Box^s p \vee \Box^s{\sim}p))$

- Sara knows that the database is consistent, but she doesn't know that it is incomplete.

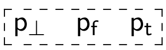

$\mathsf{p_\bot} \models \Box^s\neg(p \wedge {\sim}p) \wedge \sim\Box^s\neg(p \vee {\sim}p)$

- Sara knows whether the database is consistent.

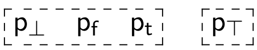

$\mathsf{p_\bot} \models \Box^s\neg(p \wedge {\sim}p) \vee \Box^s\neg\neg(p \wedge {\sim}p)$.

- Sara knows whether the database is consistent. Hendrik does not. Sara says to Hendrik: "The system manager just informed me that the database is consistent."

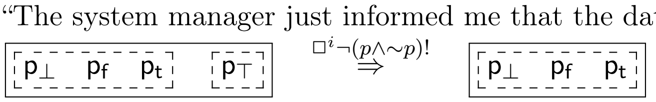

- Sara and Hendrik are both uncertain about the status of the database. The system manager says: "I will now inform Sara whether the database is consistent." He proceeds to do so, but by whispering into her ear, so that Hendrik cannot hear what he says to Sara.

  

  Here, α represents the whisper action. This is non-deterministic choice between action models α_k and α_l (and where for $\alpha_k \cup \alpha_l$ we write α), where $\alpha_k = ((K, R_\alpha, Pre_\alpha), k)$ such that $K = \{k, l\}$, $Pre_\alpha(k) = \neg(p \wedge {\sim}p)$; $Pre_\alpha(l) = \neg\neg(p \wedge {\sim}p)$; $R^s_\alpha(k,k) = \mathsf{t}$, $R^s_\alpha(l,l) = \mathsf{t}$, $R^s_\alpha(k,l) = \mathsf{f}$, and $R^s_\alpha(l,k) = \mathsf{f}$; $R^h_\alpha(i,j) = \mathsf{t}$ for all $i, j \in \{k, l\}$. Action structure α_l is the same as α_k but with l as designated point.

In all the above, we only considered a single propositional variable, p. However, we can also consider situations wherein Hendrik Edeling is the expert on p, and controls that database, whereas Sara Burgerhart (possibly) has information on the tulip's petals. Are they round and wide, or are they narrow and sleek? Let that be a proposition q. (In fact Acuminata tulips have sleek petals — they approach more the Turkish ideal tulip than the Dutch ideal tulip.) We could think of her as controlling another database. Both databases could contain thousands of items of possibly inconsistent information. The scenarios merely represent the most elementary setting to reason about database consistency and completeness by interacting agents.

Four-valued accessibility relations Our framework does not only permit four-valued propositions but also four-valued relations. Using Fitting's [66] fitting words:

> *Now, two kinds of judgments are possible. 1) A is true in situation w; and 2) w is a situation that should be considered.*

Here A is any proposition, for which we tend to write φ, and we call w a world. Fitting considered many-valued logics in general, whereas we are in bilattice logic, with judgements on truth and falsity. In other words, if $R(s, s') = \mathsf{t}$ then s' is in, and if $R(s, s') = \mathsf{f}$ then s' is out.

Let a Kripke model $\mathcal{M} = \langle S, R, V\rangle$ be given with a two-valued relation R ($(s, s') \in R$ or $(s, s') \notin R$). Let S' be a set of worlds disjoint from S. Consider $\mathcal{M}'$ with domain $S \cup S'$ and define a relation R' by $R'(s, s') = R'(s', s) = R'(s', s') = \mathsf{f}$ for any world $s' \in S'$ and any $s \in S$. Then for all φ, $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}', s) \models \varphi$. This follows easily, as f is the bottom of the truth order $\leq_t$. For non-modal formulas it is obvious that $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}', s) \models \varphi$; for modal formulas we can observe that $V(\Diamond\varphi, s) = \bigvee\{R'(s, s') * V(\varphi, s') \ : \ s' \in S\} = \bigvee\{R'(s, s') * V(\varphi, s') \ : \ s' \in S \cup S'\}$, because when $s' \in S'$ we have that $R'(s, s') * V(\varphi, s') = \mathsf{f}$. Thus, this conjunct does not affect the value of the join. Similarly, $R'(s, s') \rightarrow V(\varphi, s') = \mathsf{t}$ does not affect the value of the meet defining $V(\Box\varphi, s)$.

Not surprisingly, with values $\bot$ or $\top$ for pairs in the accessibility relation it becomes harder to appeal to our modelling intuitions. For example, what does it mean that 'Hendrik Edeling considers world s possible' has value $\top$? Does he then consider it possible and impossible at the same time? Our previous visualization with boxes is no longer suitable, and from now on we depict all pairs in the accessibility relation explicitly as arrows, labelled with the value of that pair in R (so, for example, below we have that $R(\mathsf{p_f}, \mathsf{p_t}) = \top$).

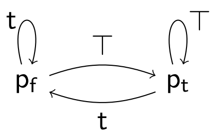

We could interpret this by saying that Hendrik's beliefs are more inclined towards p being false than towards p being true, as $\top$ is lower in the $\leq_{\mathsf{t}}$ hierarchy than t (worlds considered t are more *plausible* than worlds considered $\top$). Still, $\top$ access is good enough to get to know p. Compare the following three (distinct) models:

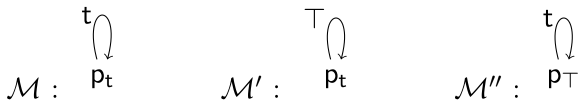

In $\mathcal{M}$ and $\mathcal{M}'$, $\Box p$ is true, whereas in $\mathcal{M}''$, $\Box p$ is false. (In $\mathcal{M}'$, $V(\Box p, \mathsf{p_t}) = R(\mathsf{p_t}, \mathsf{p_t}) \to V(p, \mathsf{p_t}) = \top \to \mathsf{t} = \mathsf{t}$; whereas in $\mathcal{M}''$, $V(\Box p, \mathsf{p}_\top) = R(\mathsf{p}_\top, \mathsf{p}_\top) \to V(p, \mathsf{p}_\top) = \mathsf{t} \to \top = \mathsf{f}$.) The latter is easily explained: $\Box p$ is false if there is an accessible world where p is false. And value $\top$ means that p is (also) false. To understand that $\Box p$ is true in $\mathcal{M}'$, it is sufficient to observe that the $\mathsf{p_t}$ world *is considered*. It is in. That it is simultaneously out does not hurt. So Hendrik still knows that tulips are red.

What properties are satisfied by Kripke models with four-valued relations that are used to interpret knowledge modalities? Are they still equivalence relations? Take transitivity: if $(s, s') \in R$ and $(s', s'') \in R$ then $(s, s'') \in R$; but if $(s, s') \notin R$ and $(s', s'') \notin R$ then we need not have that $(s, s'') \notin R$ (for example, suppose $s'' = s$). Transitivity plays a role in the four-valued logic **BS4** of [127] (employing two-valued relations), and transitivity of four-valued relations is summarily discussed in [67] in the context of combining knowledge of different experts. The answer to our questions is in the logic, not in the structures: for transitivity we need the properties enforcing the validity of $\Box\varphi \to \Box\Box\varphi$. We can achieve this with simple means. First, an example.

Hendrik Edeling knows the colour of the tulips in the Acuminata database. They are red, or white, or blue. His model of uncertainty is

$\boxed{r \quad r'}$ $\boxed{s \quad s'}$ $\boxed{b \quad b'}$

There are three equivalence classes, and all pairs are either in or out (for $(x, y) \in R$ read $R(x, y) = \mathsf{t}$ and for $(x, y) \notin R$ read $R(x, y) = \mathsf{f}$). We have that: $R(r, s) = \mathsf{f}$ and $R(s, r') = \mathsf{f}$ but $R(r, r') = \mathsf{t}$; $R(r, s) = \mathsf{f}$ and $R(s, b) = \mathsf{f}$ and $R(r, b) = \mathsf{f}$; $R(r, s) = \mathsf{f}$ and $R(s, s') = \mathsf{t}$ and $R(r, s') = \mathsf{f}$. All combinations are possible except that $R(x, y) = \mathsf{t}$ and $R(y, z) = \mathsf{t}$ imply $R(x, z) = \mathsf{t}$. That is only what matters: t or $\top$ related worlds should relate the same to all other worlds.

The structural requirements to enforce the validity of the properties of knowledge are as follows.

- If $R(s, x) = \mathsf{t}$ and $R(x, y) = \mathsf{t}$, then $R(s, y) = \mathsf{t}$.
- If $R(s, x) = \top$ and $R(x, y) = \top$, then $R(s, y) = \top$.
- If $R(s, x) = \mathsf{t}$ and $R(x, y) = \top$, then $R(s, y) = \top$.
- If $R(s, x) = \top$ and $R(x, y) = \mathsf{t}$, then $R(s, y) = \mathsf{t}$.
- $R(s, s) \in \{\mathsf{t}, \top\}$.
- If $R(s, x) \in \{\mathsf{t}, \top\}$ and $R(s, y) = i$, then $R(x, y) = i$ (where $i = \bot, \mathsf{t}, \mathsf{f}, \top$).

These cannot be properly called 'frame properties', as the manipulation of the pairs in the relation depends on their *values* in a given model. If these properties are fulfilled, then the schemata $\Box\varphi \to \varphi$, $\Box\varphi \to \Box\Box\varphi$, and $\Diamond\varphi \to \Box\Diamond\varphi$ are all valid (this is easy to see). Similarly, we get $\Box^i\varphi \to \Box^i\Box^i\varphi$, etc., for multi-agent bilattice epistemic logic.

Four-valued action models In our logical framework not only the accessibility relations of Kripke models are four-valued but also the accessibility relations of action models. Let us see some variations on the announcement of p. We have replaced the *names* of action models by their *preconditions*. (The boxes only serve to separate models and have no meaning.)

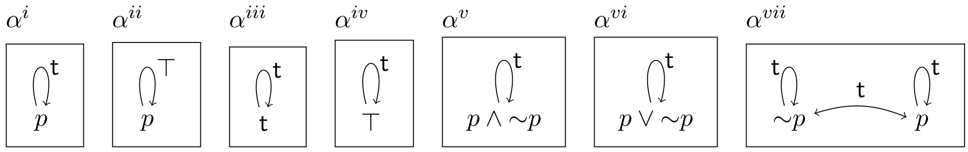

Action α_i is the public announcement of p (and also its correspondent in bilattice logic [139]). The difference between α_i and α_{ii} is that, when executed on a two-valued Kripke model, all links between worlds get value $\top$ instead of t; and in both cases the domain is restricted to the p-worlds (i.e., the $V(p) \in \{\mathsf{t}, \top\}$ worlds). For example, in a Kripke model $\mathcal{M}$ with two indistinguishable, two-valued, p and $\sim p$ worlds, both $[\alpha^i_p]\Box p$ and $[\alpha^{ii}_p]\Box p$ are true. The difference between α^i and α^{ii} only appears when evaluating knowledge of inconsistencies: $[\alpha^i_p]\Box(p \wedge \sim p)$ is false whereas $[\alpha^{ii}_p]\Box(p \wedge \sim p)$ is true, as $\mathsf{t} \to \top = \mathsf{f}$ whereas $\top \to \top = \top$. Action structures α^{iii} and α^{iv} result (when executed on a given model) in isomorphic models: they have the same update effect (namely, none at all); the worlds preserved by precondition t are the same as those preserved by precondition $\top$ (a public announcement of φ restricts the domain to worlds where $\varphi \in \{\mathsf{t}, \top\}$; trivially, $\mathsf{t} \in \{\mathsf{t}, \top\}$ and $\top \in \{\mathsf{t}, \top\}$). Actions α^v and α^{vi} we have already discussed: these are the announcements that p is inconsistent, respectively, that p is complete. Action α^{vii} has the same update effect as α^{vi}. Actions α^{vi} and α^{vii} are different from α^{iii} and α^{iv}: the latter two preserve $\bot$ worlds at their execution, the former two not. Given that, an interesting eighth version, with the same update effect as α^{iii} and α^{iv}, is:

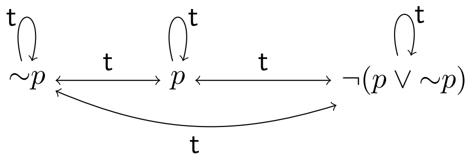

Now consider the following four alternative depictions as action models of a public announcement of φ. The rightmost of the two points (in case there are two) is the designated point.

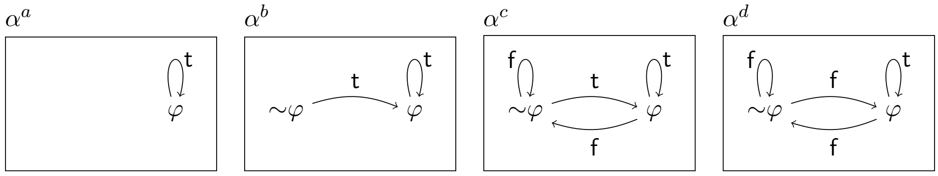

Again, α^a $(= \alpha^i)$ is the standard. Action structure α^b is known as the Gerbrandy-style *conscious update* [75]. Instead of eliminating *worlds* that do not satisfy the announcement formula, it eliminates *arrows* (pairs in the accessibility relation) that do not point to worlds satisfying the announcement. An obvious 'four-valued completion' of this action model is α^c. A less obvious four-valued completion of α^a is α^d. Clearly the update effect of α^a and α^d is the same, and also the update effect of α^b and α^c. Actions α^a and α^b have also the same update effect (where it is important that the φ-world is the designated point of α^b; the correspondence only holds when the announcement is true). This does not change for bilattice modal logic (it is about accessibility). Thus, all four describe essentially the same action!

Similarly to above we could add a third point to α^d with precondition $\neg(\sim\varphi \vee \varphi)$, which is f-accessible from and to the other points, and while keeping the φ point as the designated one.

Let this be α^e. Again, α^e has the same update effect as all the others. But executing α^e does not restrict the domain of the model. On any model, we get the same result (logically indistinguishable results) by arrow elimination when executing α^e as by world elimination when executing α^a. This can be applied to any action model: given any Kripke model $\mathcal{M}$ with domain S and action model α with domain K, once having computed the $|K|$-fold coproduct of S (cartesian product $S \times K$), we need not restrict the *domain* as when computing $\mathcal{M}_\alpha$, but it suffices to restrict the *accessibility relation*, i.e., we need to make enough $R(v, s)$ swap their value from `t` to `f`.

> **What is an announcement in four-valued logic?** *As is well known, public* announcement *logic is a misnomer, it is rather a logic of public, information changing,* events. *Various communicative phenomena including (informative) announcements count as (information changing) events: (a) an oral announcement heard by all; (b) a visual observation (by all) of a property of surrounding objects, for example, when you see a red tulip blossoming in the fields; (c) written information observed by all, such as a teacher writing* $1 + 1 = 2$ *on a blackboard, or an envelope containing information on* p*, opened in public. (Some events called 'public announcements' are not information changing events at all, but factual changing events, as in "I hereby declare Donald Trump to be the president of the USA." We exclude those from consideration.) Not all of these make sense in a setting where inconsistency or incompleteness plays a role. Direct observations are hardly ever inconsistent. A tulip is red. Or it is not red. Now it may be red or orange, or something indefinable in between. But then we would say that the proposition 'the tulip is red' is in between true and false; we would not say that it is simultaneously true and false. A visual illusion might count as a contradictory observation (*$\top$*): is the image below that of a young or of an old woman?*

> *And what would it mean that a direct observation is absent (*$\bot$*)? Whereas the contents of a letter can easily be contradictory or absent. You open it. It contains a leaf, with* p *written on it. Or the leaf contains* $\sim p$*. Or there was no leaf enclosed. Or two, one with* p *and the other with* $\sim p$*.*

4.6 Conclusions and future research

We proposed a four-valued bilattice-based modal logic including dynamic modalities for the consequences of actions. Our logic is suitable for reasoning about inconsistent and incomplete information, and about change of information in such settings. We have presented an axiomatisation of the logic and shown completeness using algebraic logic and duality theory. We hope that our logic may be useful in computer science applications.

The present paper is part of an ongoing enterprise that aims, on the applied logic side, at extending dynamic epistemic logics beyond classical reasoning and, on the theoretical side, at achieving a better understanding of the very mechanism of epistemic updates. From the latter point of view, an intriguing direction for future research is the investigation of the most general conditions for the algebraic/duality theoretic machinery to be applicable to epistemic updates. The papers [114, 105, 138, 139, 35] have shown that a uniform methodology, with few ad hoc adjustments, can be extended from the classical setting to those of intuitionistic, bilattice and finite-valued Łukasiewicz modal logics. Other logics are likely to be easily dealt with, for example positive (i.e. negation-free) modal logic and semilattice-based modal systems. The question then arises what could be minimal requirements of algebraic/relational semantics that would allow for a uniform definition of epistemic updates, perhaps one that does not heavily rely (as is so far the case) on the particular algebraic language involved. For example, since the pseudo-quotient construction involves the definition of a (partial) congruence by certain algebraic terms, we may wonder what kind of terms we should postulate in an abstract setting. Algebraic logic may turn out to be helpful here, and in particular the results from the general theory of the algebraization of logics that establish a link between logical filters (theories of a logic) and congruences of the associated algebraic semantics.

5

Neighbourhood contingency bisimulation

Contents

5.1 Introduction

A proposition is non-contingent if it is necessarily true or necessarily false, and otherwise it is contingent. The notion of (non-)contingency goes back to Aristotle [33]. The modal logic of contingency goes back to Montgomery & Routley [122]. They captured non-contingency by an operator Δ such that $\Delta\varphi$ means that formula φ is non-contingent (and where $\nabla\varphi$ means that φ is contingent). In an epistemic modal logic, 'φ is non-contingent' means that you know whether φ, and 'φ is contingent' means that you are ignorant about φ [153, 151, 63]. Contingency is definable with necessity: $\Delta\varphi$ is definable as $\Box\varphi \vee \Box\neg\varphi$. But necessity cannot always be defined with non-contingency. The definability of $\Box$ with Δ has been explored in various studies [62, 122, 134]. In [62] the *almost-definability* schema $\nabla\psi \to (\Box\varphi \leftrightarrow (\Delta\varphi \wedge \Delta(\psi \to \varphi)))$ is proposed — as long as there is a contingent proposition ψ, $\Box$ is definable with Δ; which inspired a matching notion of *contingency bisimulation*: back and forth only apply when non-bisimilar accessible worlds exist.

Schemas such as $\Delta(\varphi \wedge \psi) \to (\Delta\varphi \wedge \Delta\psi)$ are invalid for the non-contingency operator. The operator Δ is therefore not monotone, and the logic of contingency is not a normal modal logic.

Non-normal logics are standardly interpreted on neighbourhood models [37, 146, 121]. Fan & Van Ditmarsch proposed in [61] to interpret the contingency operator on neighbourhood models. They left as an open question what a suitable notion of contingency bisimulation would be over neighbourhood models. In this chapter, we answer this question.

The main contributions of this chapter are: (i) the introduction of a notion of *neighbourhood* Δ-*bisimilarity*, inspired by the semantics of the Δ-modality and [89], where different notions of structural invariance among neighbourhood models were studied using the coalgebraic representation of neighbourhood structures. By way of augmented neighbourhood models and their correspondence to Kripke models we can provide a detailed comparison to the bisimulations of [62]. We show that the two notions differ at the level of relations, but the ensuing bisimilarity notions coincide; (ii) the introduction of the notions of Δ-morphisms and Δ-quotients which leads us to prove analogues of results from [89], namely Hennessy-Milner theorem; (iii) characterisation of neighbourhood contingency logic as the Δ-bisimulation invariant fragment of classical modal logic and of first-order logic (similar to [62, Theorem 4.4, Theorem 4.5]) ; (iv) a model theoretic proof of Craig interpolation for neighbourhood contingency logic.

Overview. Section 5.2 provides preliminaries on sets, functions and relations and fixes the notations. Section 5.3 recalls contingency logic over Kripke models and introduces different perspectives on relational contingency bisimulation. Section 5.4 introduces neighbourhood contingency bisimulation, studies its properties, and provides Hennessy-Milner style theorem for an appropriate notion of saturated models. It is then followed by the characterisation results in Section 5.6. In Section 5.7 we prove Craig interpolation for neighbourhood contingency logic. The concluding Section 3.6 reflects on the relevance of our work and indicates future directions.

5.2 Preliminaries

We assume that the reader is familiar with the standard notions of sets, functions and relations. The following is merely to fix notations and to introduce the crucial notion of coherence.

5.2.1 Sets, functions and relations

Let X and U be sets. Given $U \subseteq X$, we write in_U for the inclusion map $\mathsf{in}_U \colon U \to X$, and we denote by U^c the complement of U in X. The disjoint union of two sets X_1 and X_2 is denoted by $X_1 + X_2$ and the inclusion maps by $\mathsf{in}_i \colon X_i \to X_1 + X_2$, $i = l, r$. Given a function $f \colon X \to Y$, the f-image of $U \subseteq X$ is $f[U] = \{f(x) \in Y \ : \ x \in U\}$, and the inverse f-image of $V \subseteq Y$ is $f^{-1}[V] = \{x \in X \ : \ f(x) \in V\}$. The *graph of* f is the relation $Gr(f) = \{(x, f(x)) \subseteq X \times Y : x \in X\}$. The *kernel of* f is the relation $\ker(f) = \{(x, y) \in X \times X \ : \ f(x) = f(y)\}$. Let $R \subseteq X \times Y$ be a relation. We denote by $\pi_l : R \to X$ and $\pi_r : R \to Y$ the left and the right projections of R, respectively. The R-image of $U \subseteq X$ is the set $R[U] = \{y \in Y \ : \ \exists x \in U : (x, y) \in R\}$, and the inverse R-image of $V \subseteq Y$ is $R^{-1}[V] = \{x \in X \ : \ \exists y \in V : (x, y) \in R\}$.

Given a relation $R \subseteq X \times Y$, the relation $R^{-1} = \{(y, x) \in Y \times X \ : \ (x, y) \in R\}$ is the *converse* of R, and the *composition* of R and the relation $S \subseteq Y \times Z$ is written as $R; S$ and defined by $R; S = \{(x, z) \in X \times Z : \exists y \in Y$ such that $(x, y) \in R$ and $(y, z) \in S\}$. Let $R \subseteq X \times X$. The reflexive closure of R is $R^r = R \cup \mathrm{Id}_X$, where $\mathrm{Id}_X = \{(x, x) \ : \ x \in X\}$. The symmetric closure of R is $R^s = R \cup R^{-1}$. (Note that R is symmetric if $R = R^{-1}$.) The transitive closure of R is $R^+ = \bigcup_{i \geq 1} R^i$, where R^i is defined inductively by $R^1 = R$ and $R^{i+1} = R^i; R$ for $i \geq 1$. The equivalence closure of R is defined as $R^e = ((R^r)^s)^+$. If R is an equivalence relation, we often write $[x]_R$ (or simply $[x]$) instead of $R(x)$. Every relation $R \subseteq X \times Y$ can be viewed as a relation on $X + Y$ which is defined as $R_{X+Y} = \{(\mathsf{in}_l(x), \mathsf{in}_l(y)) : (x, y) \in R\}$.

5.2.2 Coherence

Throughout this chapter the notion of *coherence* will be used extensively.

Definition 5.2.1 (R-coherent pairs) *Let $R \subseteq X \times Y$ be a relation, $U \subseteq X$ and $V \subseteq Y$. The pair (U, V) is R-*coherent *if $R[U] \subseteq V$ and $R^{-1}[V] \subseteq U$, or equivalently, for all $(x, y) \in R$, $x \in U$ iff $y \in V$. Given a relation $R \subseteq X \times X$, we say that $U \subseteq X$ is R-*closed *if (U, U) is R-coherent.* $\dashv$

For every relation $R \subseteq X \times Y$, trivially $(\emptyset, \emptyset)$ and (X, Y) are R-coherent. Let us consider a simple example.

Example 5.2.2 *Suppose $X = \{1, 2, 3\}$, $Y = \{a, b, c\}$. Let $R = \{(1, a), (2, a), (3, b)\}$ be a relation between X and Y. Then, the following are all R-coherent pairs: $(\emptyset, \emptyset)$, $(\emptyset, \{c\})$, $(\{1, 2\}, \{a\})$, $(\{1, 2\}, \{a, c\})$, $(\{1, 2, 3\}, \{a, b\})$, $(\{1, 2, 3\}, \{a, b, c\})$.* $\dashv$

We list a number of useful (easily proven) properties of R-coherence in the following lemmas.

Lemma 5.2.3 *Let $R \subseteq X \times Y$ a relation, and let $U \subseteq X$ and $V \subseteq Y$.*

1. *If $R \subseteq R'$, where $R' \subseteq X \times Y$, and (U, V) is R'-coherent, then (U, V) is R-coherent.*
2. *If $R = Gr(f)$ for a function $f \colon X \to Y$, and (U, V) is R-coherent, then $U = f^{-1}[V]$.*
3. *(U, V) is R-coherent iff $(U + V, U + V)$ is R_{X+Y}-coherent.*
4. *(U, V) is R-coherent iff $\pi_l^{-1}[U] = \pi_r^{-1}[V]$.* $\dashv$

Lemma 5.2.4 *Let $R \subseteq X \times X$ be a relation and $U \subseteq X$. Then, U is R^e-closed iff U is a union of R^e-equivalence classes.* $\dashv$

5.3 Contingency Logic

In this section we recall the syntax and Kripke semantics of basic modal logic and contingency logic, and contingency bisimulation following [62, 63]. We introduce a novel notion of relational contingency bisimulation defined in terms of coherence, and compare it to the contingency bisimulation of [62].

Definition 5.3.1 (Languages) *Let At be a set of atomic propositions. The languages $\mathcal{L}_\Box$ and $\mathcal{L}_\Delta$ are generated by the following grammars:*

$$\mathcal{L}_\Box \ni \varphi \quad ::= \quad p \in \mathsf{At} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi$$
$$\mathcal{L}_\Delta \ni \varphi \quad ::= \quad p \in \mathsf{At} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Delta\varphi$$

The other Boolean connectives f, t, $\vee$ and $\leftrightarrow$ are defined in the usual way.

The formula $\Delta\varphi$ should be read as 'φ is non-contingent', that is, φ is necessarily true or φ is necessarily false. The language $\mathcal{L}_\Delta$ can be viewed as a fragment of $\mathcal{L}_\Box$ via an inductively defined translation $(-)^t : \mathcal{L}_\Delta \to \mathcal{L}_\Box$ with only non-trivial clause $(\Delta\varphi)^t = \Box\varphi^t \vee \Box\neg\varphi^t$.

We interpret the language $\mathcal{L}_\Delta$ on Kripke models (Def. 3.2.2) as follows:

Definition 5.3.2 (Semantics) *Let $\mathcal{M} = (S, R, V)$ be a Kripke model, and $s \in S$. The interpretation of formulas from $\mathcal{L}_\Delta$ is defined inductively in the usual manner:*

$$
\begin{array}{lll}
(\mathcal{M}, s) \models p & \textit{iff} & s \in V(p) \\
(\mathcal{M}, s) \models \varphi \wedge \psi & \textit{iff} & (\mathcal{M}, s) \models \varphi \textit{ and } (\mathcal{M}, s) \models \psi \\
(\mathcal{M}, s) \models \neg\varphi & \textit{iff} & (\mathcal{M}, s) \not\models \varphi \\
(\mathcal{M}, s) \models \Delta\varphi & \textit{iff} & \textit{for all } t_1, t_2 \in R(s) : (\mathcal{M}, t_1 \models \varphi \Leftrightarrow (\mathcal{M}, t_2) \models \varphi).
\end{array}
$$

where $p \in \mathsf{At}$. We say that $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ are (modally) $\mathcal{L}_\Delta$-equivalent *(notation: $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}', s')$) if for all $\varphi \in \mathcal{L}_\Delta$, $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}', s') \models \varphi$.* ⊣

We say $\varphi \in \mathcal{L}_\Delta$ is *valid* in a Kripke model $\mathcal{M} = \langle S, R, V\rangle$, denoted by $\mathcal{M} \models \varphi$, if $(\mathcal{M}, s) \models \varphi$ for all $s \in S$, and it is *valid*, denoted by $\models_{\mathbf{CL}} \varphi$, if $\mathcal{M} \models \varphi$ for every Kripke model. These definitions can be easily extended to sets of formulas in the following way: a set $\Phi \subseteq \mathcal{L}_\Delta$ is valid in a Kripke model $\mathcal{M} = \langle S, R, V\rangle$, if $\mathcal{M} \models \varphi$ for all $\varphi \in \Phi$. Let $\Phi \cup \{\varphi\} \subseteq \mathcal{L}_\Delta$. We write $\Phi \models \varphi$, if φ is a *local semantic consequence* of Φ, if for any model $\mathcal{M}$ and s in $\mathcal{M}$, if $(\mathcal{M}, s) \models \Phi$, then $(\mathcal{M}, s) \models \varphi$. We say Φ is *consistent* if $\Phi \not\models \mathsf{f}$, which means there is a Kripke model that satisfies Φ. Finally, we define *contingency logic* **CL** as the set of validities over the class of Kripke models in the language $\mathcal{L}_\Delta$, that is, for all formulas $\varphi \in \mathcal{L}_\Delta$: $\varphi \in \mathbf{CL}$ iff $\models_{\mathbf{CL}} \varphi$.

For all pointed Kripke models $(\mathcal{M}, s)$, and all $\varphi \in \mathcal{L}_\Delta$, $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}, s) \models \varphi^t$.

Recall the definition of Kripke bisimulations from Chapter 3 (see Def. 3.2.6, page 40). In [62], Fan, Wang & Van Ditmarsch defined a weaker notion (for Δ) which we refer to as *o-Δ-bisimulation* for "original Δ-bisimulation".

Definition 5.3.3 (o-Δ-bisimulation) *Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model. A relation $Z \subseteq S \times S$ is an* o-Δ-bisimulation *on $\mathcal{M}$, if whenever $(s, s') \in Z$:*

(Atoms) *s and s' satisfy the same propositional variables;*
(Δ-Zig) *for all $t \in R(s)$, if there are $t_1, t_2 \in R(s)$ such that $(t_1, t_2) \notin Z$, then there is a $t' \in R(s')$ such that $(t, t') \in Z$;*
(Δ-Zag) *for all $t' \in R(s')$, if there are $t'_1, t'_2 \in R(s')$ such that $(t'_1, t'_2) \notin Z$, then there is a $t \in R(s)$ such that $(t, t') \in Z$.*

We write $(\mathcal{M}, s) \approx^{\mathsf{on}}_\Delta (\mathcal{M}, s')$, if there is an o-$\Delta$-bisimulation on $\mathcal{M}$ that contains (s, s'). We say two pointed Kripke models $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ are o-Δ-bisimilar*, written $(\mathcal{M}, s) \approx_\Delta (\mathcal{M}', s')$, if $(\mathcal{M} + \mathcal{M}', \mathsf{in}_l(s)) \approx^{\mathsf{on}}_\Delta (\mathcal{M} + \mathcal{M}', \mathsf{in}_r(s'))$, i.e., there is an o-$\Delta$-bisimulation on the disjoint union of $\mathcal{M}$ and $\mathcal{M}'$ linking (the injection images of) s and s', where the disjoint union of two Kripke models is the disjoint union of $\mathcal{M}$ and $\mathcal{M}'$ is the structure $\mathcal{M} + \mathcal{M}' = \langle S^+, R^+, V^+\rangle$ in which $S^+ = S \biguplus S'$, $R^+ = R \biguplus R'$, and $V^+(p) = V(p) \biguplus V'(p)$ for every $p \in \mathsf{At}$.* ⊣

Note that $(\mathcal{M}, s) \approx_\Delta (\mathcal{M}', s')$ is not witnessed by a relation $Z \subseteq S \times S'$ since, by definition, o-Δ-bisimulation relations always live on a single model.

We introduced the notation $\approx^{\mathsf{on}}_\Delta$, since, a priori, it is not clear whether $(\mathcal{M}, s) \approx^{\mathsf{on}}_\Delta (\mathcal{M}, s')$ iff $(\mathcal{M}, s) \approx_\Delta (\mathcal{M}, s')$. The next proposition shows that this is true, and hence we could dispense with the notation $\approx^{\mathsf{on}}_\Delta$, but for now we keep writing $\approx^{\mathsf{on}}_\Delta$ for clarity.

Proposition 5.3.4 *Let $\mathcal{M}$ be a Kripke model that contains two states s and s'. Then we have:*

$$(\mathcal{M}, s) \approx^{\mathsf{on}}_\Delta (\mathcal{M}, s') \quad \textit{iff} \quad (\mathcal{M}, s) \approx_\Delta (\mathcal{M}, s').$$

Proof ($\Rightarrow$): If Z is an o-Δ-bisimulation on $\mathcal{M}$, then it is easy to prove that $Y := \{(\mathsf{in}_l(s), \mathsf{in}_r(t)) : (s, t) \in Z\}$ is an o-Δ-bisimulation on $\mathcal{M} + \mathcal{M}$.

($\Leftarrow$): Let Y be an o-Δ-bisimulation on $\mathcal{M}+\mathcal{M}$. We denote by R_l the accessibility relation of the left component of $\mathcal{M}+\mathcal{M}$, and by R_r the accessibility relation of the right component of $\mathcal{M}+\mathcal{M}$. Define $Z := \{(s,s') \in S \times S \ : \ \exists i,j \in \{l,r\} : (\mathsf{in}_i(s), \mathsf{in}_j(s')) \in Y\}$. To prove Δ-Zig for Z, suppose $(s,s') \in Z$ and $t,t_1,t_2 \in R(s)$ such that $(t_1,t_2) \notin Z$. This implies that $\mathsf{in}_i(t), \mathsf{in}_i(t_1), \mathsf{in}_i(t_2) \in R_i(\mathsf{in}_i(s))$, and there are $i,j \in \{l,r\}$ such that $(\mathsf{in}_i(s), \mathsf{in}_j(s')) \in Y$, and by definition of Z, $(\mathsf{in}_i(t_1), \mathsf{in}_i(t_2)) \notin Y$. By Δ-Zig for Y, there are $\mathsf{in}_j(t'), \mathsf{in}_j(t'_1), \mathsf{in}_j(t'_2) \in R_j(\mathsf{in}_j(s'))$ such that $(\mathsf{in}_i(t), \mathsf{in}_j(t')), (\mathsf{in}_i(t_1), \mathsf{in}_j(t'_1)), (\mathsf{in}_i(t_2), \mathsf{in}_j(t'_2)) \in Y$. Hence $t' \in R(s')$ and $(t,t'), (t_1,t'_1), (t_2,t'_2) \in Z$, which proves Δ-Zig. The condition Δ-Zag can be proved in a similar manner. □

Given a model $\mathcal{M}$, we will also view $\approx_\Delta^{\mathsf{on}}$ as the relation on the state space of $\mathcal{M}$ that contains all pairs (s,s') such that $(\mathcal{M},s) \approx_\Delta^{\mathsf{on}} (\mathcal{M},s')$. In order to compare o-$\Delta$-bisimilarity with our later notion (in Definition 5.3.6), we need the following result.

Proposition 5.3.5 *On each Kripke model, $\approx_\Delta^{\mathsf{on}}$ is an o-Δ-bisimulation that is an equivalence relation.* ⊣

Proof We will first show that the set of o-Δ-bisimulations on a Kripke model $\mathcal{M}$ is closed under taking unions, converse, and transitive symmetric closure. Let $\mathcal{M} = \langle S,R,V \rangle$ be a Kripke model. Closure under unions and converse can be straightforwardly proved. To show closure under transitive symmetric closure, let X be an o-Δ-bisimulation on $\mathcal{M}$, and let Z be the symmetric closure of X. As the converse of an o-Δ-bisimulation is an o-Δ-bisimulation, it follows that Z is an o-Δ-bisimulation. We show Z^+ is an o-Δ-bisimulation, as well. Suppose $(s,s') \in Z^+$. The clause (**Atoms**) is trivial. To prove **(Δ-Zig)**, let $t,t_1,t_2 \in R(s)$ be such that $(t_1,t_2) \notin Z^+$ (i.e., for all $n \geq 1$, $(t_1,t_2) \notin Z^n$). To produce a $t' \in R(s')$ such that $(t,t') \in Z^+$, we claim that for all $m \geq 1$, for all $u,v \in S$, if $(s,u) \in Z^m$ and $v \in R(s)$ then there is a $t' \in R(u)$ such that $(v,t') \in Z^m$, where s is from our main assumption. This suffices to prove **(Δ-Zig)** for Z^+, since $(s,s') \in Z^+$ implies that there is an $m \geq 1$ such that $(s,s') \in Z^m$. From $t \in R(s)$, we get from the claim a $t' \in R(s')$ such that $(t,t') \in Z^m$, hence also $(t,t') \in Z^+$. A similar argument shows that **(Δ-Zag)** also holds for Z^+. Now we prove the claim by induction on m.

Base case ($m=1$): Let $(s,u) \in Z$ and $v \in R(s)$. From $(t_1,t_2) \notin Z^+$ follows $(t_1,t_2) \notin Z$. From **(Δ-Zig)** for Z we obtain a $t' \in R(u)$ such that $(v,t') \in Z$.

Induction step ($m+1$): Let $(s,u) \in Z^{m+1}$ and $v \in R(s)$. Then there is a u' such that $(s,u') \in Z^m$ and $(u',u) \in Z$. Applying the induction hypothesis to $(s,u') \in Z^m$ and $v,t_1,t_2 \in R(s)$ (where t_1 and t_2 are from our main assumption), we obtain $v',v'_1,v'_2 \in R(u')$ such that $(v,v'), (t_1,v'_1), (t_2,v'_2) \in Z^m$. Towards a contradiction, suppose that $(v'_1,v'_2) \in Z$. Then from $(t_1,v'_1), (t_2,v'_2) \in Z^m$, $(v'_1,v'_2) \in Z$ and symmetry of Z it follows that $(t_1,t_2) \in Z^{2m+1} \subseteq Z^+$. This contradicts the main assumption that $(t_1,t_2) \notin Z^+$. Therefore, $(v'_1,v'_2) \notin Z$. Together with $v',v'_1,v'_2 \in R(u')$, we now get from **(Δ-Zig)** for Z a $t' \in R(u)$ such that $(v',t') \in Z$ and $(v,t') \in Z^{m+1}$, as required.

Now using the above closure properties, we will show that $\approx_\Delta^{\mathsf{on}}$ is an o-Δ-bisimulation that is an equivalence relation. By definition, the relation $\approx_\Delta^{\mathsf{on}}$ on $\mathcal{M}$ is the union of all o-Δ-bisimulations on $\mathcal{M}$, and hence the largest one. Reflexivity of $\approx_\Delta^{\mathsf{on}}$ is immediate since the identity relation is an o-Δ-bisimulation. Symmetry follows from closure under converse. For transitivity, we use that the composition of two bisimulations is contained in the transitive symmetric closure of their union, which is again a bisimulation. □

We note that the above proposition follows from the stronger result that says o-Δ-bisimilarity is an equivalence relation over the class of all pointed Kripke models [62]. Unlike the case for standard Kripke bisimulation, this is quite non-trivial to prove, since o-Δ-bisimulations are not closed under composition (Example 5.3.7).

It is shown in [62] that Kripke bisimilarity implies o-Δ-bisimilarity, but the following example shows that the other direction does not hold.

$$\mathcal{M} \quad s : p \longrightarrow s_1 : p \qquad\qquad \mathcal{M}' \quad t : p$$

The pointed models $(\mathcal{M}, s)$ and $(\mathcal{M}', t)$ are o-Δ-bisimilar, but there is no Kripke bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ that links the states s and t.

In the following we will recall some known results about *o*-Δ-bisimilarity from [62]. The notion of o-Δ-bisimilarity is *adequate* for $\mathcal{L}_\Delta$, that is, o-Δ-bisimilarity implies $\mathcal{L}_\Delta$-equivalence [62, Proposition 3.5], whereas the converse only holds for $\mathcal{L}_\Delta$-*saturated* [25] Kripke models [62, Proposition 3.9]. Given a Kripke model $\mathcal{M} = \langle S, R, V\rangle$, we say $\mathcal{M}$ is $\mathcal{L}_\Delta$-saturated if for every $\Gamma \subseteq \mathcal{L}_\Delta$ and every $s \in S$ we have: if every *finite* subset of Γ is satisfiable in $R(s)$ then Γ is satisfiable in $R(s)$. Contingency logic has been characterised as the o-Δ-bisimulation invariant fragment within basic modal logic and within first order logic; that is, an $\mathcal{L}_\Box$-formula (first-order formula) is equivalent to an $\mathcal{L}_\Delta$-formula iff it is invariant under o-Δ-bisimulation [62, Theorem 4.4, 4.5].

The notion of contingency bisimulation for neighbourhood models using coherent sets, introduced later in Definition 5.4.4, has a natural analogue for Kripke models. The main intuition behind this definition is the semantics of the Δ-modality over Kripke models. Given a Kripke model $\mathcal{M} = \langle S, R, V\rangle$ and $s \in S$, the satisfaction condition for the formulas of the form $\Delta\varphi$ is

$$(\mathcal{M}, s) \models \Delta\varphi \quad \text{iff} \quad \text{for all } t_1, t_2 \in R(s) : (\mathcal{M}, t_1 \models \varphi \Leftrightarrow (\mathcal{M}, t_2) \models \varphi)$$

an it can be equivalently written as follows:

$$(\mathcal{M}, s) \models \Delta\varphi \quad \text{iff} \quad R(s) \subseteq [\![\varphi]\!]_\mathcal{M} \quad \text{or} \quad R(s) \subseteq ([\![\varphi]\!]_\mathcal{M})^c. \tag{5.1}$$

where $[\![\varphi]\!]_\mathcal{M} = \{s \in S : (\mathcal{M}, s) \models \varphi\}$ denoted the truth set of φ in $\mathcal{M}$. Condition (5.1) gives rise to the following definition.

Definition 5.3.6 (rel-Δ-bisimulation) *Let $\mathcal{M} = \langle S, R, V\rangle$ and $\mathcal{M}' = \langle S', R', V'\rangle$ be Kripke models. A relation $Z \subseteq S \times S'$ is a* rel-Δ-bisimulation *(for* relational Δ-bisimulation*) between $\mathcal{M}$ and $\mathcal{M}'$, if whenever $(s, s') \in Z$:*

(Atoms) *s and s' satisfy the same propositional variables;*
(Coherence) *for all Z-coherent pairs (U, U'):*

$$(R(s) \subseteq U \text{ or } R(s) \subseteq U^c) \quad \textit{iff} \quad (R'(s') \subseteq U' \text{ or } R'(s') \subseteq U'^c)$$

We write $(\mathcal{M}, s) \sim^{\mathsf{betw}}_\Delta (\mathcal{M}', s')$, when there is a rel-$\Delta$-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ that contains (s, s'). A rel-Δ-bisimulation on a model $\mathcal{M}$ *is a rel-Δ-bisimulation between $\mathcal{M}$ and $\mathcal{M}$. We define the notion of rel-Δ-bisimilarity between states in potentially different models via the disjoint union (analogously to the notion of o-Δ-bisimilarity): We say two pointed models $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ are* rel-Δ-bisimilar*, written $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s')$, if $(\mathcal{M} + \mathcal{M}', \mathsf{in}_l(s)) \sim^{\mathsf{betw}}_\Delta (\mathcal{M} + \mathcal{M}', \mathsf{in}_r(s'))$, i.e,. if there is a rel-Δ-bisimulation on $\mathcal{M} + \mathcal{M}'$ that contains* $(\mathsf{in}_l(s), \mathsf{in}_r(s'))$.⊣

In Proposition 5.3.12 we will see that over a single model, $\sim^{\mathsf{betw}}_\Delta$ and $\sim_\Delta$ coincide, but in general they differ. At first it would seem more natural to define rel-Δ-bisimilarity between pointed models as $\sim^{\mathsf{betw}}_\Delta$. However, the following Example 5.3.7 (item 4) shows that this notion is too restrictive. The example also shows that, in general, rel-Δ-bisimulations are different from o-Δ-bisimulations.

Example 5.3.7 *Consider the four figures (and matching items below) where we assume a single variable p to be false in all states of all models, except in figure* 4 *where p is true at s and t.*

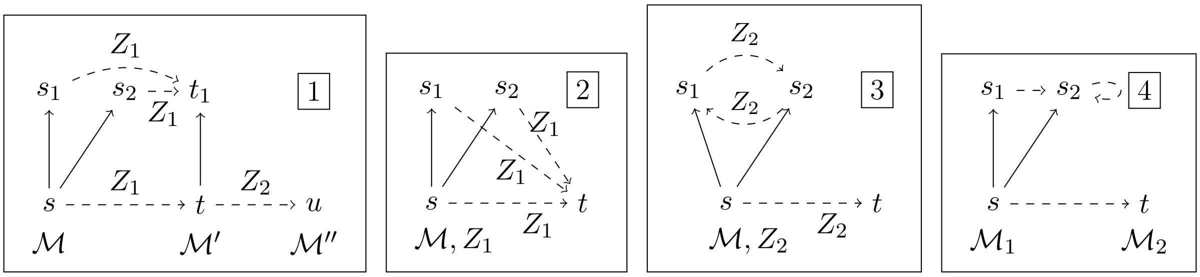

1 The composition of two o-Δ-bisimulations may not be an o-Δ-bisimulation. *The relations Z_1 and Z_2 are o-Δ-bisimulations, but not $Z_1; Z_2 = \{(s, u)\}$, since $s_1, s_2 \in R(s)$ and there is no successor of u that is $Z_1; Z_2$-related to s_1 and s_2. Hence, Z_1 fails to satisfy the Δ-Zig condition.*

2 A rel-Δ-bisimulation may not be an o-Δ-bisimulation. *The relation Z_1 is not an o-Δ-bisimulation, as Δ-Zig fails for $(s,t) \in Z_1$ because $s_1, s_2 \in R(s)$ and there is no successor of t that is Z_1-related to s_1 and s_2. However, Z_1 is a rel-Δ-bisimulation on $\mathcal{M}$. The Z_1-coherent pairs are: $(\{s, s_1, s_2\}, U')$ and (S, U') for all U' with $t \in U'$, $(\{t\}, U')$ for all U' with $t \notin U'$, and $(\emptyset, \emptyset)$. Since $R(s_1) = R(s_2) = R(t) = \emptyset$,* ***(Coherence)*** *for (s_1, t) and (s_2, t) is satisfied. For (s, t), e.g., for $(\{s, s_1, s_2\}, \{t\})$: $R(s) = \{s_1, s_2\} \subseteq \{s, s_1, s_2\}$ and $R(t) = \emptyset \subseteq \{t\}$, and for $(\{t\}, \{s_1\})$: $R(s) = \{s_1, s_2\} \subseteq \{t\}^c$ and $R(t) = \emptyset \subseteq \{s_1\}$.*

3 An o-Δ-bisimulation may not be a rel-Δ-bisimulation. *It is easy to check that the relation Z_2 is an o-Δ-bisimulation, but not a rel-Δ-bisimulation, since $(\{s_1\}, \{s_2\})$ is Z_2-coherent, $(s,t) \in Z_2$, and $\emptyset = R(t) \subseteq \{s_2\}$, but $R(s) \nsubseteq \{s_1\}$ and $R(s) \nsubseteq \{s_1\}^c$. Therefore, Z_2 is not a rel-Δ-bisimulation.*

4 A rel-Δ-bisimulation on a disjoint union, but not between disjoints. *The pictured relation is a rel-Δ-bisimulation on $\mathcal{M}_1 + \mathcal{M}_2$, but there is no rel-Δ-bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2$ linking s and t. Since p is only true in s and t and not in s_1 and s_2, the only candidate is $\{(s,t)\}$, but the coherent pair $(\{s, s_1\}, \{t\})$ does not satisfy* **(Coherence)***, since $R_1(s) \nsubseteq \{s, s_1\}$, whereas $R_2(t) = \emptyset \subseteq \{t\}$. Hence, $(\mathcal{M}_1 + \mathcal{M}_2, \mathsf{in}_l(s)) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}_1 + \mathcal{M}_2, \mathsf{in}_r(t))$, but not $(\mathcal{M}_1, s) \sim_\Delta (\mathcal{M}_2, t)$.* ⊣

Although the two notions of contingency bisimulations differ at the level of relations, we can show that rel-Δ-bisimilarity coincides with o-Δ-bisimilarity. We will need the following lemma.

Lemma 5.3.8 *Let $\mathcal{M} = \langle S, R, V\rangle$ be a Kripke model, and assume that $Z \subseteq S \times S$ is an equivalence relation. Z is an o-Δ-bisimulation iff Z is a rel-Δ-bisimulation.* ⊣

Proof First, suppose Z is an o-Δ-bisimulation and $(s, s') \in Z$. Since Z is an equivalence relation, we need to show that for all Z-closed subsets U,

$$(R(s) \subseteq U \text{ or } R(s) \subseteq U^c) \quad \text{iff} \quad (R(s') \subseteq U \text{ or } R(s') \subseteq U^c) \tag{5.2}$$

To see that (5.2) holds, let $R(s) \subseteq U$ *or* $R(s) \subseteq U^c$, where U is Z-closed. Suppose towards a contradiction that $R(s') \cap U \neq \emptyset$ and $R(s') \cap U^c \neq \emptyset$. Then, there are $t_1, t_2 \in R(s')$ such that $t_1 \in U$ and $t_2 \in U^c$. Since U is Z-closed, $(t_1, t_2) \notin Z$. By applying Δ-Zag, there are $s_1, s_2 \in R(s)$ such that $(s_1, t_1), (s_2, t_2) \in Z$. From $R(s) \subseteq U$ *or* $R(s) \subseteq U^c$, we obtain $s_1, s_2 \in U$ or $s_1, s_2 \in U^c$

and since U is Z-closed, it follows that $t_1, t_2 \in U$ or $t_1, t_2 \in U^c$, which is a contradiction. Therefore, $R(s') \subseteq U$ or $R(s') \subseteq U^c$. The other direction of (5.2) may be checked in a similar way.

Now, assume that Z is a rel-Δ-bisimulation, and let $(s, s') \in Z$. (**Atoms**) is immediate. For **(Δ-Zig)**, assume $t, t_1, t_2 \in R(s)$ such that $(t_1, t_2) \notin Z$. Suppose towards a contradiction that there is no $t' \in R(s')$ such that $(t, t') \in Z$, then $Z(t) \cap R(s') = \emptyset$ and hence $R(s') \subseteq (Z(t))^c$. As $Z(t)$ is Z-closed and Z is a rel-Δ-bisimulation we get by **(Coherence)** that $R(s) \subseteq Z(t)$ or $R(s) \subseteq (Z(t))^c$. But since Z is an equivalence relation and $(t_1, t_2) \notin Z$, it implies that $R(s) \subseteq Z(t)$ is false. On the other hand, $R(s) \subseteq (Z(t))^c$ is also false since Z is an equivalence relation and $t \in R(s) \cap Z(t)$. Hence we have a contradiction and conclude that Z satisfies the **(Δ-Zig)** condition. By a similar argument Z satisfies **(Δ-Zag)**. □

We have an analogue of Proposition 5.3.5 for rel-Δ-bisimilarity (Proposition 5.3.10), namely the relation $\sim_\Delta^{\mathsf{betw}}$ is the largest rel-Δ-bisimulation on a given Kripke model $\mathcal{M}$, and it is an equivalence relation. It can be proved in a similar way as in the proof of Proposition 5.3.5 via the following closure properties.

Lemma 5.3.9 *Let $\mathcal{M} = \langle S, R, V\rangle$ and $\mathcal{M}' = \langle S', R', V'\rangle$ be Kripke models.*

1. *The identity relation $\mathrm{Id}_S \subseteq S \times S$ is a rel-Δ-bisimulation on $\mathcal{M}$.*
2. *If $Z \subseteq S \times S'$ is a rel-Δ-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ then $Z^{-1} \subseteq S' \times S$ is a rel-Δ-bisimulation between $\mathcal{M}'$ and $\mathcal{M}$.*
3. *The set of rel-Δ-bisimulations are closed under arbitrary unions: If $Z_i \subseteq S \times S'$, $i \in I$, are rel-Δ-bisimulations, then so is $\bigcup_{i\in I} Z_i$.*
4. *If $Z \subseteq S \times S$ a rel-Δ-bisimulation on $\mathcal{M}$, then the reflexive, symmetric closure of Z is also a rel-Δ-bisimulation.*
5. *If $Z \subseteq S \times S$ is a reflexive, symmetric rel-Δ-bisimulation on $\mathcal{M}$, then the transitive closure Z^+ is also a rel-Δ-bisimulation.* ⊣

Proof Items 1-3 are straightforward, details are left to the reader. Item 4 follows from the previous three, since the reflexive closure of a relation $Z \subseteq S \times S$ is $Z \cup \mathrm{Id}_S$, and the symmetric closure is $Z \cup Z^{-1}$.

To prove item 5, we show that for all $n \geq 1$, Z^n is a rel-Δ-bisimulation. It follows by closure under unions that $Z^+ = \bigcup_{n\geq 1} Z^n$ is a rel-Δ-bisimulation. The proof is by induction on n.

The base case $(n = 1)$ holds by assumption on Z. Assume it holds for n.

Induction step $(n+1)$: First note that if (U, U'') is Z^{n+1}-coherent, then since Z^{n+1} is reflexive, it follows that $U = U''$. Moreover, since Z and Z^n are reflexive and $Z^{n+1} = Z^n; Z$, it follows that $Z \subseteq Z^{n+1}$ and $Z^n \subseteq Z^{n+1}$. As (U, U) is Z^{n+1}-coherent, we obtain that (U, U) is Z-coherent as well as Z^n-coherent. Now suppose $(s, s') \in Z^n$, $(s', s'') \in Z$ and (U, U) is Z^{n+1}-coherent. We then have

$$\begin{array}{rcll} R(s) \subseteq U \text{ or } R(s) \subseteq U^c & \iff & R(s') \subseteq U \text{ or } R(s') \subseteq U^c & \text{(by ind.hyp.)} \\ & \iff & R(s'') \subseteq U \text{ or } R(s'') \subseteq U^c & (Z \text{ is rel-}\Delta\text{-bisim.}) \end{array}$$

Hence, Z^{n+1} is a rel-Δ-bisimulation. □

Proposition 5.3.10 *On each Kripke model, $\sim_\Delta^{\mathsf{betw}}$ is a rel-Δ-bisimulation that is an equivalence relation.* ⊣

Proof By definition, $\sim^{\mathsf{betw}}_{\Delta}$ is the union of all rel-Δ-bisimulations on $\mathcal{M}$, which is again a rel-Δ-bisimulation due to closure under unions (Lemma 5.3.9(3)), and hence the largest one. Also, the equivalence closure of a rel-Δ-bisimulation on $\mathcal{M}$ is again one (Lemma 5.3.9(5)). Hence $\sim^{\mathsf{betw}}_{\Delta}$ on a model $\mathcal{M}$ is the largest rel-Δ-bisimulation on $\mathcal{M}$ and an equivalence relation. □

It follows from Propositions 5.3.5 and 5.3.10, and Lemma 5.3.8 that the two notions of contingency bisimilarity coincide.

Proposition 5.3.11 *Let $\mathcal{M}$ and $\mathcal{M}'$ be Kripke models.*

1. *For all s, t in $\mathcal{M}$: $(\mathcal{M}, s) \approx^{\mathsf{on}}_{\Delta} (\mathcal{M}, t)$ iff $(\mathcal{M}, s) \sim^{\mathsf{betw}}_{\Delta} (\mathcal{M}, t)$.*
2. *For all s in $\mathcal{M}$ and s' in $\mathcal{M}'$: $(\mathcal{M}, s) \approx_{\Delta} (\mathcal{M}', s')$ iff $(\mathcal{M}, s) \sim_{\Delta} (\mathcal{M}', s')$.* ⊣

Proof *Item 1:* Let $\mathcal{M}$ be a Kripke model. From Proposition 5.3.5, the relation $\approx^{\mathsf{on}}_{\Delta}$ on $\mathcal{M}$ is an equivalence relation and an o-Δ-bisimulation on $\mathcal{M}$. From Proposition 5.3.10, the relation $\sim^{\mathsf{betw}}_{\Delta}$ on $\mathcal{M}$ is an equivalence relation and a rel-Δ-bisimulation on $\mathcal{M}$. The result now follows from Lemma 5.3.8. *Item 2:* Follows from the definition of $\approx_{\Delta}$ and $\sim_{\Delta}$ over pointed models, and item 1. □

Recall that [59, 60] proved that over the class of all pointed Kripke models, o-Δ-bisimilarity $\approx_{\Delta}$ is an equivalence. Due to Proposition 5.3.11(2), it follows that also rel-Δ-bisimilarity $\sim_{\Delta}$ an equivalence.

Finally, we show that we can dispense with the notation $\approx^{\mathsf{on}}_{\Delta}$ as item 1 of the next proposition ensures that no ambiguity can arise when writing $(\mathcal{M}, s) \approx_{\Delta} (\mathcal{M}, s')$. We also clarify the similar question regarding $\sim^{\mathsf{betw}}_{\Delta}$ and $\sim_{\Delta}$.

Proposition 5.3.12 *For all Kripke models $\mathcal{M}$ and $\mathcal{M}'$:*

1. *$(\mathcal{M}, s) \approx^{\mathsf{on}}_{\Delta} (\mathcal{M}, s')$ iff $(\mathcal{M}, s) \approx_{\Delta} (\mathcal{M}, s')$.*
2. *$(\mathcal{M}, s) \sim^{\mathsf{betw}}_{\Delta} (\mathcal{M}', s')$ implies $(\mathcal{M}, s) \sim_{\Delta} (\mathcal{M}', s')$. The implication is strict.*
3. *$(\mathcal{M}, s) \sim^{\mathsf{betw}}_{\Delta} (\mathcal{M}, s')$ iff $(\mathcal{M}, s) \sim_{\Delta} (\mathcal{M}, s')$.* ⊣

Proof *Item 1.* ($\Rightarrow$): If Z is an o-Δ-bisimulation on $\mathcal{M}$, then it is easy to prove that $Y := \{(\mathsf{in}_l(s), \mathsf{in}_r(t)) : (s,t) \in Z\}$ is an o-Δ-bisimulation on $\mathcal{M} + \mathcal{M}$.

($\Leftarrow$): Let Y be an o-Δ-bisimulation on $\mathcal{M} + \mathcal{M}$. We denote by R_l the accessibility relation of the left component of $\mathcal{M} + \mathcal{M}$, and by R_r the accessibility relation of the right component of $\mathcal{M} + \mathcal{M}$. Define $Z := \{(s, s') \in S \times S : \exists i, j \in \{l, r\} : (\mathsf{in}_i(s), \mathsf{in}_j(s')) \in Y\}$. To prove Δ-Zig for Z, suppose $(s, s') \in Z$ and $t, t_1, t_2 \in R(s)$ such that $(t_1, t_2) \notin Z$. This implies that $\mathsf{in}_i(t), \mathsf{in}_i(t_1), \mathsf{in}_i(t_2) \in R_i(\mathsf{in}_i(s))$, and there are $i, j \in \{l, r\}$ such that $(\mathsf{in}_i(s), \mathsf{in}_j(s')) \in Y$, and by definition of Z, $(\mathsf{in}_i(t_1), \mathsf{in}_i(t_2)) \notin Y$. By Δ-Zig for Y, there are $\mathsf{in}_j(t'), \mathsf{in}_j(t'_1), \mathsf{in}_j(t'_2) \in R_j(\mathsf{in}_j(s'))$ such that $(\mathsf{in}_i(t), \mathsf{in}_j(t')), (\mathsf{in}_i(t_1), \mathsf{in}_j(t'_1)), (\mathsf{in}_i(t_2), \mathsf{in}_j(t'_2)) \in Y$. Hence $t' \in R(s')$ and $(t, t'), (t_1, t'_1), (t_2, t'_2) \in Z$, which proves Δ-Zig. The condition Δ-Zag can be proved in a similar manner.

Item 2.
$$
\begin{aligned}
(\mathcal{M}, s) \sim^{\mathsf{betw}}_{\Delta} (\mathcal{M}, s') &\iff (\mathcal{M}, s) \approx^{\mathsf{on}}_{\Delta} (\mathcal{M}, s') && \text{Proposition 5.3.11(2)}\\
&\iff (\mathcal{M}, s) \approx_{\Delta} (\mathcal{M}, s') && \text{(Item 1)}\\
&\iff (\mathcal{M}, s) \sim_{\Delta} (\mathcal{M}, s') && \text{Proposition 5.3.11(1)}
\end{aligned}
$$

Item 3. The implication can be proved using Lemma 5.4.6 and Proposition 5.4.5 of the next section. The converse fails since item 4 of Example 5.3.7 shows models $(\mathcal{M}_1, s)$ and $(\mathcal{M}_2, t)$ such that $(\mathcal{M}_1, s) \sim_{\Delta} (\mathcal{M}_2, t)$, however, we do not have $(\mathcal{M}_1, s) \sim^{\mathsf{betw}}_{\Delta} (\mathcal{M}_2, t)$. □

5.4 Neighbourhood Semantics of Contingency Logic

In this section we recall the neighbourhood semantics of $\mathcal{L}_\Delta$ from [63], and then we proceed to introduce the notion of Δ-bisimulation between neighbourhood models, and investigate its properties.

Definition 5.4.1 (Neighbourhood models) *A neighbourhood frame is a pair (S, ν) where S is a set of states and $\nu : S \to \mathcal{P}(\mathcal{P}(S))$ is a neighbourhood function which assigns to each $s \in S$ its collection $\nu(s)$ of neighbourhoods. A neighbourhood model is a triple $\mathcal{M} = (S, \nu, V)$ where (S, ν) is a neighbourhood frame and $V : \mathsf{At} \to \mathcal{P}(S)$ is a valuation. A* neighbourhood morphism *between $\mathcal{M} = (S, \nu, V)$ and $\mathcal{M}' = (S', \nu', V')$ is a function $f : S \to S'$ such that (i) for all $p \in \mathsf{At}$, $s \in V(p)$ iff $f(s) \in V'(p)$, and (ii) for all subsets $U \subseteq S'$, $f^{-1}[U] \in \nu(s)$ iff $U \in \nu'(f(s))$.* ⊣

Neighbourhood morphisms are the neighbourhood analogue of bounded morphisms, and they indeed preserve truth of $\mathcal{L}_\Box$-formulas [89, Lem. 2.6], and hence also of $\mathcal{L}_\Delta$-formulas. The semantics of $\mathcal{L}_\Box \cup \mathcal{L}_\Delta$-formulas is given below in Definition 5.4.3.

In what follows, we will also use disjoint unions (or coproducts) of neighbourhood models. We recall the definition from [89, Def. 2.9].

Definition 5.4.2 (Disjoint union of neighbourhood models) *Let $\mathcal{M}_1 = (S_1, \nu_1, V_1)$ and $\mathcal{M}_2 = (S_2, \nu_2, V_2)$ be two neighbourhood models. Their* disjoint union *$\mathcal{M}_1 + \mathcal{M}_2$ is the model $\mathcal{M} = (S, \nu, V)$ where $S = S_1 + S_2$, $V(p) = \mathsf{in}_l[V_1(p)] \cup \mathsf{in}_r[V_2(p)]$, and for all $U \subseteq S_1 + S_2$, all $i = l, r$, and all $s_i \in S_i$:*

$$U \in \nu(\mathsf{in}_i(s_i)) \quad \Longleftrightarrow \quad \mathsf{in}_i^{-1}[S_i] \in \nu_i(s_i) \tag{5.3}$$

Being a bit sloppy and omitting explicit use of inclusion maps, the condition (5.3) can be stated as: $U \in \nu(s_i)$ iff $U \cap S_i \in \nu_i(s_i)$. The definition of ν ensures that the inclusion maps $\mathsf{in}_l : S_1 \to S_1 + S_2$ and $\mathsf{in}_r : S_2 \to S_1 + S_2$ are neighbourhood morphisms, and hence preserve truth of $\mathcal{L}_\Box \cup \mathcal{L}_\Delta$-formulas.

Definition 5.4.3 (Neighbourhood Semantics of Contingency Logic) *Given a neighbourhood model $\mathcal{M} = (S, \nu, V)$. The interpretation of formulas from $\mathcal{L}_\Box$ and $\mathcal{L}_\Delta$ in $\mathcal{M}$ is defined inductively for atomic propositions and Boolean connectives as usual. Truth of modal formulas is given by,*

$(\mathcal{M}, s) \models \Box\varphi$ *iff* $[\![\varphi]\!]_\mathcal{M} \in \nu(s)$
$(\mathcal{M}, s) \models \Delta\varphi$ *iff* $[\![\varphi]\!]_\mathcal{M} \in \nu(s)$ *or* $[\![\varphi]\!]^c_\mathcal{M} \in \nu(s)$.

*where $[\![\varphi]\!]_\mathcal{M} = \{s \in S : (\mathcal{M}, s) \models \varphi\}$ denotes the truth set of φ in $\mathcal{M}$. We write $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}', s')$ if $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ satisfy the same $\mathcal{L}_\Delta$-formulas. A subset $X \subseteq S$ is $\mathcal{L}_\Delta$-*definable*, if there is a $\varphi \in \mathcal{L}_\Delta$ such that $X = [\![\varphi]\!]_\mathcal{M}$.*

Truth and validity of a formula φ of $\mathcal{L}_\Delta$ (or $\mathcal{L}_\Box$) over neighbourhood models are defined in the same way as for Kripke semantics. Here we apply similar notational conventions as we have set in Section 5.3. The set of $\mathcal{L}_\Box$-validities over neighbourhood models is known as *classical modal logic* [37] , and we denote it by **ML**. We define *classical contingency logic* **CCL** to be the set of $\mathcal{L}_\Delta$-validities over the class of neighbourhood models.

Again, it is clear that over neighbourhood models we can view $\mathcal{L}_\Delta$ as a fragment of $\mathcal{L}_\Box$, since for all neighbourhood models $\mathcal{M}$, all states s in $\mathcal{M}$, and all $\varphi \in \mathcal{L}_\Delta$, $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}, s) \models \varphi^t$ (page 106).

We now recall how Kripke models can be viewed as certain neighbourhood models [37]. For a Kripke model $\mathcal{M} = (S, R, V)$, define $nbh(\mathcal{M}) = (S, \nu_R, V)$ where $\nu_R(s) = \{X \subseteq S : R(s) \subseteq X\}$. It is straightforward to check that for all $\varphi \in \mathcal{L}_\Box \cup \mathcal{L}_\Delta$,

$$(\mathcal{M}, s) \models \varphi \quad \text{iff} \quad (nbh(\mathcal{M}), s) \models \varphi. \tag{5.4}$$

A neighbourhood model (S, ν, V) is *augmented* (cf. [37]) if all neighbourhood collections are closed under supersets and under arbitrary intersections, that is, for all $s \in S$, if $U \in \nu(s)$ and $U \subseteq U' \subseteq S$, then $U' \in \nu(s)$; and $\bigcap \nu_R(s) \in \nu_R(s)$. For an augmented $\mathcal{M} = (S, \nu, V)$, define a Kripke model $krp(\mathcal{M}) = (S, R, V)$ by taking $R(s) = \bigcap \nu(s)$. Again, $\mathcal{M}$ and $krp(\mathcal{M})$ are pointwise equivalent, and we have $nbh(krp(\mathcal{M})) = \mathcal{M}$ and $krp(nbh(K)) = K$. Thus, Kripke models are in 1-1 correspondence with augmented neighbourhood models.

A proof system $\mathbb{CL}$ for contingency logic (**CL**) has been proposed by Fan et al. in [63] and they showed that $\mathbb{CL}$ is sound and strongly complete with respect to the class of Kripke frames [63, Theorem 19]. From equation (5.4) it follows immediately that $\mathbb{CL}$ is sound and strongly complete with respect to the class of augmented neighbourhood frames. This question was left open in [61].

We now define the notion of Δ-bisimulation between neighbourhood models. The definition is inspired by the notion of *precocongruences* [89, Def. 5.3.3] and the semantics of the Δ-modality. To see this, let us recall the satisfaction condition of Δ-formulas over neighbourhood models. Given a neighbourhood model $\mathcal{M}$ and state s in $\mathcal{M}$,

$$(\mathcal{M}, s) \models \Delta\varphi \quad \text{iff} \quad [\![\varphi]\!]_{\mathcal{M}} \in \nu(s) \;\text{ or }\; [\![\varphi]\!]^c_{\mathcal{M}} \in \nu(s). \tag{5.5}$$

Now, our notion of bisimulation is derived from the formulation of (5.5) using the notion of coherent pairs.

Definition 5.4.4 (nbh-Δ-bisimulation) *Let $\mathcal{M} = (S, \nu, V)$ and $\mathcal{M}' = (S', \nu', V')$ be neighbourhood models. A relation $Z \subseteq S \times S'$ is a* nbh-Δ-bisimulation *(for "neighbourhood Δ-bisimulation") if for all $(s, s') \in Z$, the following hold:*

(Atoms) *s and s' satisfy the same atomic propositions.*

(Coherence) *for all Z-coherent pairs (U, U'):*

$$U \in \nu(s) \text{ or } U^c \in \nu(s) \quad \text{iff} \quad U' \in \nu'(s') \text{ or } U'^c \in \nu'(s').$$

We write $(\mathcal{M}, s) \sim^{\mathsf{betw}}_\Delta (\mathcal{M}', s')$, if there is a nbh-$\Delta$-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ that contains (s, s'). A nbh-Δ-bisimulation on a model $\mathcal{M}$ *is a nbh-Δ-bisimulation between $\mathcal{M}$ and $\mathcal{M}$. For the same reason that we need to define the notion of rel-Δ-bisimilarity on disjoint union of two Kripke models, we also need to define the notion of nbhd-Δ-bisimilarity between two potentially different models via disjoint union. We say two pointed models $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ are* nbh-Δ-bisimilar, *written $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s')$, if $(\mathcal{M} + \mathcal{M}', \mathsf{in}_l(s)) \sim^{\mathsf{betw}}_\Delta (\mathcal{M} + \mathcal{M}', \mathsf{in}_r(s'))$, i.e,. if there is a nbh-Δ-bisimulation on $\mathcal{M} + \mathcal{M}'$ that contains $(\mathsf{in}_l(s), \mathsf{in}_r(s'))$.* ⊣

The following proposition shows that there is no conflict between the notions of nbh-Δ-bisimulations and rel-Δ-bisimulations for augmented models. This allows us to simply speak of Δ-bisimulations, and it justifies the overloading of the notation $\sim_\Delta$.

Proposition 5.4.5 *A relation Z is a rel-Δ-bisimulation between Kripke models $\mathcal{M}$ and $\mathcal{M}'$ if and only if Z is a nbh-Δ-bisimulation between $nbh(\mathcal{M})$ and $nbh(\mathcal{M}')$. Consequently,*

1. $(\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}', s')$ *iff* $(nbh(\mathcal{M}), s) \sim_\Delta^{\mathsf{betw}} (nbh(\mathcal{M}'), s')$.
2. $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s')$ *iff* $(nbh(\mathcal{M}), s) \sim_\Delta (nbh(\mathcal{M}'), s')$. ⊣

Proof Item 1 is straightforward to prove using the correspondence between Kripke models and augmented neighbourhood models. Let $\mathcal{M} = (S, R, V)$ and $\mathcal{M}' = (S', R', V')$ be two Kripke models. By definition of ν_R and $\nu_{R'}$, we have:

$$\begin{array}{llcll} (a) & U \in \nu_R(s) \text{ or } U^c \in \nu_R(s) & \iff & (b) & R(s) \subseteq U \text{ or } R(s) \subseteq U^c \qquad \text{and} \\ (a') & U' \in \nu_{R'}(s') \text{ or } U'^c \in \nu_{R'}(s') & \iff & (b') & R'(s') \subseteq U' \text{ or } R'(s') \subseteq U'^c \end{array}$$

If Z is a rel-Δ-bisimulation, we have $(b) \iff (b')$, and it follows that Z is a nbh-Δ-bisimulation. Conversely, if Z is a nbh-Δ-bisimulation, then we have $(a) \iff (a')$, and hence Z is a rel-Δ-bisimulation. Hence it is clear that item 1 holds.

For item 2 we use item 1 and the isomorphism $nbh(\mathcal{M} + \mathcal{M}') \cong nbh(\mathcal{M}) + nbh(\mathcal{M}')$, which is easy to verify. We have

$$\begin{array}{lcl} (\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s') & \iff & \text{(def.)} \\ (\mathcal{M} + \mathcal{M}', s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M} + \mathcal{M}', s') & \iff & \text{(item 1)} \\ (nbh(\mathcal{M} + \mathcal{M}'), s) \sim_\Delta^{\mathsf{betw}} (nbh(\mathcal{M} + \mathcal{M}'), s') & \iff & \text{(isomorphism)} \\ (nbh(\mathcal{M}) + nbh(\mathcal{M}'), s) \sim_\Delta^{\mathsf{betw}} (nbh(\mathcal{M}) + nbh(\mathcal{M}'), s') & \iff & \text{(def.)} \\ (nbh(\mathcal{M}), s) \sim_\Delta (nbh(\mathcal{M}'), s') & & \end{array}$$

□

Over arbitrary pointed neighbourhood models, $\sim_\Delta^{\mathsf{betw}}$ is strictly contained in $\sim_\Delta$, but on a single neighbourhood model they coincide.

Lemma 5.4.6 *For all pointed neighbourhood models $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$:*

1. $(\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}', s')$ *implies* $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s')$. *The implication is strict.*
2. $(\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}, s')$ *iff* $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}, s')$. ⊣

Proof *Item 1.* Let $\mathcal{M}_1 = (S_1, \nu_1, V_1)$ and $\mathcal{M}_2 = (S_2, \nu_2, V_2)$ be neighbourhood models. We show that if Z is a nbh-Δ-bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2$, then the embedding $\mathsf{in}(Z) = \{(\mathsf{in}_l(s_1), \mathsf{in}_r(s_2)) : (s_1, s_2) \in Z\}$ is a nbh-Δ-bisimulation on $\mathcal{M}_1 + \mathcal{M}_2 = (S, \nu, V)$. We will use the following fact about complements. For all $X_1 \subseteq S_1$ and $X_2 \subseteq S_2$:

$$(S_1 + S_2) \setminus (X_1 + X_2) = (S_1 \setminus X_1) + (S_2 \setminus X_2).$$

Now let $(\mathsf{in}_l(s_1), \mathsf{in}_r(s_2)) \in \mathsf{in}(Z)$ and let $(U_1 + U_2, U_1' + U_2')$ be $\mathsf{in}(Z)$-coherent where $U_1, U_1' \subseteq S_1$ and $U_2, U_2' \subseteq S_2$. As $(U_1 + U_2, U_1' + U_2')$ is $\mathsf{in}(Z)$-coherent, it follows that (U_1, U_2') is Z-coherent, and we have:

$$\begin{array}{lcl} U_1 + U_2 \in \nu(\mathsf{in}_l(s_1)) \text{ or } (S_1 + S_2) \setminus (U_1 + U_2) \in \nu(\mathsf{in}_l(s_1)) & \iff & \text{(def. } \nu) \\ U_1 \in \nu_1(s_1) \text{ or } (S_1 \setminus U_1) \in \nu_1(s_1) & \iff & (Z \text{ is } \Delta\text{-bis.}) \\ U_2' \in \nu_2(s_2) \text{ or } (S_2 \setminus U_2') \in \nu_2(s_2) & \iff & \text{(def. } \nu) \\ U_1' + U_2' \in \nu(\mathsf{in}_r(s_2)) \text{ or } (S_1 + S_2) \setminus (U_1' + U_2') \in \nu(\mathsf{in}_r(s_2)) & & \end{array}$$

Hence, $\mathsf{in}(Z)$ is a Δ-bisimulation. The implication is strict due to Example 5.3.7 (item 4) and Proposition 5.4.5.

Item 2. ($\Rightarrow$) follows from item 1. To prove ($\Leftarrow$), assume that Z' is a nbh-Δ-bisimulation on $\mathcal{M}+\mathcal{M}$. We show that $Z := \{(s,t) \in S \times S \ : \ \exists i,j \in \{l,r\} : (\mathsf{in}_i(s), \mathsf{in}_j(t)) \in Z'\}$ is a nbh-Δ-bisimulation on $\mathcal{M}$. First, note that for all $s \in S$, $U \subseteq S$, and all $i \in \{l,r\}$: $\mathsf{in}_i(s) \in \mathsf{in}_l[U] \cup \mathsf{in}_r[U]$ iff $s \in U$.

(Atoms): Let $(s,t) \in Z$ witnessed by $(\mathsf{in}_i(s), \mathsf{in}_j(t)) \in Y$ where $i,j \in \{l,r\}$. Since Y satisfies **(Atoms)**, we have $\mathsf{in}_i(s) \in \mathsf{in}_r[V(p)] \cup \mathsf{in}_l[V(p)]$ iff $\mathsf{in}_j(t) \in \mathsf{in}_l[V(p)] \cup \mathsf{in}_r[V(p)]$, and hence $s \in V(p)$ iff $t \in V(p)$.

(Coherence): We first note that if the pair (U,V) is Z-coherent, then $(\mathsf{in}_l[U] \cup \mathsf{in}_r[U], \mathsf{in}_l[V] \cup \mathsf{in}_r[V])$ is Z'-coherent. Namely, take any pair $(\mathsf{in}_i(s), \mathsf{in}_j(t)) \in Z'$ where $i,j \in \{l,r\}$. By definition of Z, it follows that $(s,t) \in Z$. We now have $\mathsf{in}_i(s) \in \mathsf{in}_l[U] \cup \mathsf{in}_r[U]$ iff $s \in U$ iff (by Z-coherence) $t \in V$ iff $\mathsf{in}_j(t) \in \mathsf{in}_l[V] \cup \mathsf{in}_r[V]$. Furthermore, it is straightforward to show that for all $s \in S$, all $U \subseteq S$, and all $i \in \{l,r\}$:

$$U \in \nu(s) \iff (\mathsf{in}_l[U] \cup \mathsf{in}_r[U]) \in \nu'(\mathsf{in}_i(s)) \tag{5.6}$$

$$U^c \in \nu(s) \iff (\mathsf{in}_l[U] \cup \mathsf{in}_r[U])^c \in \nu'(\mathsf{in}_i(s)) \tag{5.7}$$

(Coherence) for Z now follows easily from (5.6), (5.7) and coherence for Z'. □

We state another basic fact about Δ-bisimilarity which can be proved using closure properties as for Proposition 5.3.5.

Proposition 5.4.7 *On each neighbourhood model, $\sim_\Delta$ is a Δ-bisimulation that is an equivalence relation.* ⊣

Proof It is straightforward to prove an analogue of Lemma 5.3.9 for neighbourhood models, and the proposition follows. □

As desired, Δ-bisimilar states cannot be distinguished with the $\mathcal{L}_\Delta$-language.

Proposition 5.4.8 *For all pointed neighbourhood models $(\mathcal{M}_1, s_1)$ and $(\mathcal{M}_2, s_2)$, we have*

$$(\mathcal{M}_1, s_1) \sim_\Delta (\mathcal{M}_2, s_2) \quad \textit{implies} \quad (\mathcal{M}_1, s_1) \equiv_\Delta (\mathcal{M}_2, s_2).$$

Proof By definition, $(\mathcal{M}_1, s_1) \sim_\Delta (\mathcal{M}_2, s_2)$ iff $(\mathcal{M}_1 + \mathcal{M}_2, \mathsf{in}_l(s_1)) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}_1 + \mathcal{M}_2, \mathsf{in}_r(s_2))$. Since the inclusion morphisms preserve truth, we have for all $\mathcal{L}_\Delta$-formulas φ that $(\mathcal{M}_1, s_1) \models \varphi$ iff $((\mathcal{M}_1 + \mathcal{M}_2), \mathsf{in}_l(s_1)) \models \varphi$, and similarly for $(\mathcal{M}_2, s_2)$. Hence it suffices to prove that for all models $\mathcal{M}$, $(\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}, s')$ implies $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, s')$.

So assume that Z is a Δ-bisimulation on a model $\mathcal{M}$. We prove that for all formulas $\varphi \in \mathcal{L}_\Delta$ and all $(s,s') \in Z$, $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}, s') \models \varphi$, by induction on φ. The base case $\varphi = p$ holds by (**Atoms**). The Boolean cases are routine, so let's turn to the case where $\varphi = \Delta\psi$. By induction hypothesis, we have for all $(x,y) \in Z$, $x \in [\![\psi]\!]_\mathcal{M}$ iff $y \in [\![\psi]\!]_\mathcal{M}$. That is, the pair $([\![\psi]\!]_\mathcal{M}, [\![\psi]\!]_\mathcal{M})$ is Z-coherent. As Z is a Δ-bisimulation, it follows that for all $(s,s') \in Z$, $([\![\psi]\!]_\mathcal{M} \in \nu(s)$ or $[\![\psi]\!]^c_\mathcal{M} \in \nu(s))$ iff $([\![\psi]\!]_\mathcal{M} \in \nu'(s')$ or $[\![\psi]\!]^c_\mathcal{M} \in \nu'(s'))$, that is, $(\mathcal{M}, s) \models \Delta\psi$ iff $(\mathcal{M}, s') \models \Delta\psi$. □

As with the standard notions of Kripke and neighbourhood bisimulations, $\mathcal{L}_\Delta$-equivalence does not always imply Δ-bisimilarity. Namely, by Propositions 5.3.11(2) and 5.4.5(2), Δ-bisimilarity coincides with o-Δ-bisimilarity over Kripke frames, and in [62, Example 3.10] it was shown that $\mathcal{L}_\Delta$-equivalence does not imply o-Δ-bisimilarity, hence $\mathcal{L}_\Delta$-equivalence also does not imply Δ-bisimilarity. However, a converse to Proposition 5.4.8 can be proved for an appropriate notion of saturated models following a similar line of reasoning as in [89, section 4.1]. To this end, we introduce Δ-morphisms and Δ-congruences. They will play the part of neighbourhood morphisms and congruences from [89].

Definition 5.4.9 (Δ-morphisms and Δ-congruences) *Let $\mathcal{M} = (S, \nu, V)$ and $\mathcal{M}' = (S', \nu', V')$ be neighbourhood models. A function $f\colon S \to S'$ is a* Δ-morphism from $\mathcal{M}$ to $\mathcal{M}'$ *if its graph $Gr(f)$ is a Δ-bisimulation. A relation is a* Δ-congruence *if it is the kernel of some Δ-morphism.* ⊣

It is natural to ask whether Δ-morphisms are a generalisation of neighbourhood morphisms (cf. Definition 5.4.1). This is indeed the case.

Lemma 5.4.10 *Every neighbourhood morphism is a Δ-morphism.* ⊣

Proof Let $\mathcal{M} = (S, \nu, V)$ and $\mathcal{M}' = (S', \nu', V')$ be neighbourhood models, and let $f\colon S \to S'$ be a neighbourhood morphism between them. We show that $Gr(f)$ is a Δ-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$. Let (U, U') be $Gr(f)$-coherent. Then $U = f^{-1}[U']$ and since f is a neighbourhood morphism and $f^{-1}[U']^c = f^{-1}[U'^c]$, we have:

$$\begin{array}{lcl} f^{-1}[U'] \in \nu(s) & \iff & U' \in \nu'(f(s)) \quad \text{and} \\ f^{-1}[U']^c \in \nu(s) & \iff & U'^c \in \nu'(f(s)) \end{array}$$

It follows that $Gr(f)$ is a Δ-bisimulation. □

As a step towards showing that Δ-congruences are Δ-bisimulations, we show that we can take quotients with respect to Δ-bisimulations that are also equivalence relations.

Proposition 5.4.11 (Δ-quotient) *Let $\mathcal{M} = (S, \nu, V)$ be a neighbourhood model and let Z be a Δ-bisimulation on $\mathcal{M}$ which is also an equivalence relation, i.e., for all Z-closed $U \subseteq S$ and all $(s, t) \in Z$,*

$$(U \in \nu(s) \text{ or } U^c \in \nu(s)) \iff (U \in \nu(t) \text{ or } U^c \in \nu(t)). \qquad (\dagger)$$

We define the Δ-quotient of $\mathcal{M}$ by Z *as the model $\mathcal{M}_Z = (S_Z, \nu_Z, V_Z)$ where $S_Z = \{[s] : s \in S\}$ is the set of Z-equivalence classes, $V_Z(p) = \{[s] : s \in V(p)\}$, and*

$$\nu_Z([s]) = \{U_Z \subseteq S_Z : q^{-1}[U_Z] \in \nu(s) \text{ or } q^{-1}[U_Z]^c \in \nu(s)\}.$$

The quotient map $q\colon S \to S_Z$ given by $q(s) = [s]$ is a Δ-morphism, and $Z = \ker(q)$. Consequently, $(\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}_Z, [s])$. ⊣

Proof First, we verify that ν_Z is well defined, i.e., that $[s] = [t]$ implies $\nu_Z([s]) = \nu_Z([t])$. For $\overline{U} \subseteq S_Z$, we have that $U := q^{-1}[\overline{U}] \subseteq S$ is Z-closed. Hence for all $s, t \in S$ such that $[s] = [t]$ (i.e. $(s, t) \in Z$) and all $\overline{U} \subseteq S_Z$, we have:

$$\begin{array}{lcl} \overline{U} \in \nu_Z([s]) & \iff & q^{-1}[\overline{U}] \in \nu(s) \text{ or } q^{-1}[\overline{U}]^c \in \nu(s) \\ & \iff & q^{-1}[\overline{U}] \in \nu(t) \text{ or } q^{-1}[\overline{U}]^c \in \nu(t) \\ & \iff & \overline{U} \in \nu_Z([t]) \end{array}$$

In order to check that q is a Δ-morphism, note that by Lemma 5.2.3(2) if $(U, \overline{V})$ is $Gr(q)$-coherent then $U = q^{-1}[\overline{V}]$. By the definition of ν_Z, we now have for all $s \in S$,

$$U \in \nu(s) \text{ or } U^c \in \nu(s) \quad \iff \quad \overline{V} \in \nu_Z([s])$$

Moreover, for all $\overline{V} \subseteq S_Z$, $q^{-1}[\overline{V}^c] = q^{-1}[\overline{V}]^c$, hence $\overline{V} \in \nu_Z([s])$ iff $\overline{V}^c \in \nu_Z([s])$, and we can conclude that q is a Δ-morphism. □

We can now show that Δ-congruences are indeed a special kind of Δ-bisimulations. This will be used to prove the Hennessy-Milner theorem in a moment.

Proposition 5.4.12 *Let $\mathcal{M} = (S, \nu, V)$ be a neighbourhood model and Z a relation on S. The relation Z is a Δ-congruence iff Z is an equivalence relation and a Δ-bisimulation.* $\dashv$

Proof Assume $Z = \ker(f)$ for some Δ-morphism f from $\mathcal{M}$ to $\mathcal{M}'$. Note that if U is Z-closed then $(U, f[U])$ is $Gr(f)$-coherent. Equation (†) now easily follows from f being a Δ-morphism, and $Z = \ker(f)$. Conversely, if Z is an equivalence relation and a Δ-bisimulation on $\mathcal{M}$, then we can form the Δ-quotient $\mathcal{M}_Z$, and it follows that Z is a Δ-congruence. $\square$

Proposition 5.4.12 allows us to show a neighbourhood analogue of the fact that Kripke bisimilarity implies o-Δ-bisimilarity [62]. For neighbourhood models, the equivalence notion that matches the expressiveness of the language $\mathcal{L}_\Box$ is called behavioural equivalence [89]: Two pointed neighbourhood models $(\mathcal{M}, s)$ and $(\mathcal{M}', s')$ are *behaviourally equivalent* if there exists a neighbourhood model $\mathcal{M}''$ and neighbourhood morphisms $f\colon \mathcal{M} \to \mathcal{M}''$ and $f'\colon \mathcal{M}' \to \mathcal{M}''$ such that $f(s) = f'(s')$.

Proposition 5.4.13 *Let $\mathcal{M}$ be a neighbourhood model, and s, t two states in $\mathcal{M}$. If $(\mathcal{M}, s)$ and $(\mathcal{M}, t)$ are behaviourally equivalent then they are Δ-bisimilar.* $\dashv$

Proof If $(\mathcal{M}, s)$ and $(\mathcal{M}, t)$ are behaviourally equivalent, then by [89, Prop. 3.20] the pair (s, t) is contained in a congruence, i.e. in the kernel of a neighbourhood morphism f. By Lemma 5.4.10, $\ker(f)$ is a Δ-congruence, which by Proposition 5.4.12, is a Δ-bisimulation on $\mathcal{M}$, hence $(\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}, t)$. Finally, it follows from Lemma 5.4.6 that $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}, t)$. $\square$

Finally, we prove a Hennessy-Milner style theorem for an appropriate notion of saturated models which essentially comes from [89, section 4.1].

Definition 5.4.14 ($\mathcal{L}_\Delta$-saturated model) *Let $\mathcal{M} = (S, \nu, V)$ be a neighbourhood model. A subset $X \subseteq S$ is* $\mathcal{L}_\Delta$-compact *if for all sets Φ of $\mathcal{L}_\Delta$-formulas, if any finite subset $\Phi' \subseteq \Phi$ is satisfiable in X, then Φ is satisfiable in X. $\mathcal{M}$ is* $\mathcal{L}_\Delta$-saturated, *if for all $s \in S$ and all $\equiv_\Delta$-closed neighbourhoods $X \in \nu(s)$, both X and X^c are $\mathcal{L}_\Delta$-compact.* $\dashv$

The next lemma is needed to prove the Hennessy-Milner theorem.

Lemma 5.4.15 *Let $\mathcal{M} = (S, \nu, V)$ be a neighbourhood model.*

1. *If for all $s \in S$ and all $\equiv_\Delta$-coherent neighbourhoods $X \in \nu(s)$, there is a $\varphi \in \mathcal{L}_\Delta$ such that $X = [\![\varphi]\!]_\mathcal{M}$, then $\equiv_\Delta$ is a Δ-congruence.*

2. *If $\mathcal{M}$ is $\mathcal{L}_\Delta$-saturated then for all $X \subseteq S$, X is $\equiv_\Delta$-coherent iff X is $\mathcal{L}_\Delta$-definable.*

3. *If $\mathcal{M}$ is $\mathcal{L}_\Delta$-saturated, then $\equiv_\Delta$ is Δ-congruence on $\mathcal{M}$.* $\dashv$

Proof The proof is analogous to the proof of Lemma 4.3, Lemma 4.5 of [89].
Item (1). By Prop. 5.4.12, it suffices to show that $\equiv_\Delta$ is a Δ-bisimulation. Let $s, s' \in S$ be such that $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, s')$. The condition **(Atoms)** is immediate. As for **(Coherence)**, let $s, s' \in S$ be such that $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, s')$ and let $U \subseteq S$ be $\equiv_\Delta$-coherent. By the assumption there is a $\varphi \in \mathcal{L}_\Delta$ such that $U = [\![\varphi]\!]_\mathcal{M}$. Then the semantics of $\Delta\varphi$ implies that $U \in \nu(s)$ or $U^c \in \nu(s)$ iff $(\mathcal{M}, s) \models \Delta\varphi$. But as $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, s')$, it follows that $(\mathcal{M}, s) \models \Delta\varphi$ iff $(\mathcal{M}, s') \models \Delta\varphi$.

Again by the semantics of $\Delta\varphi$ we obtain that $(\mathcal{M}, s') \models \Delta\varphi$ iff $U \in \nu(s')$ or $U^c \in \nu(s')$. This implies that $U \in \nu(s)$ or $U^c \in \nu(s)$ iff $U \in \nu(s')$ or $U^c \in \nu(s')$. Hence, $\equiv_\Delta$ is a Δ-bisimulation.

Item (2). Let $X \subseteq S$ be an $\equiv_\Delta$-coherent subset. By Lemma 5.2.4, $X = \bigcup_{i\in I}[x_i]_{\equiv_\Delta}$ for some index set I. For each $i \in I$ and $(\mathcal{M}, y) \not\equiv_\Delta (\mathcal{M}, x_i)$, there is a modal $\mathcal{L}_\Delta$-formula $\gamma_{i,y}$ such that $x_i \models \gamma_{i,y}$ and $y \models \neg\gamma_{i,y}$. So by taking $\Gamma_i = \{\gamma_{i,y} \ : \ (\mathcal{M}, y) \not\equiv_\Delta (\mathcal{M}, x_i)\}$, we have $[x_i]_{\equiv_\Delta} = \bigcap_{\gamma_{i,y}\in\Gamma_i} [\![\gamma_{i,y}]\!]$ for each $i \in I$. Since Γ_i is not satisfiable in X^c and $\mathcal{M}$ is $\mathcal{L}_\Delta$-saturated, it follows that there is ${\Gamma_i}^0 \subseteq \Gamma_i$ such that $[x]_{\equiv_\Delta} \subseteq \bigcap_{\gamma_{i,y}\in\Gamma_i^0} [\![\gamma_{i,y}]\!]$. Defining $\gamma_i = \bigwedge \Gamma_i^0$ for each $i \in I$, we therefore have $X = \bigcup_{i\in I} [\![\gamma_i]\!]_{\mathcal{M}}$. Now by $\mathcal{L}_\Delta$-compactness of X, we obtain a finite subset $\Gamma_0 \subseteq \{\gamma_i \ : \ i \in I\}$ such that $X = [\![\bigvee \Gamma_0]\!]_{\mathcal{M}}$. That is, X is definable by the formula $\gamma = \bigvee \Gamma_0$.

Item (3). It immediately follows from items 1 and 2. □

We now proceed with the Hennessy-Milner theorem.

Theorem 5.4.16 (Hennessy-Milner)

1. *For all $\mathcal{L}_\Delta$-saturated neighbourhood models $\mathcal{M}$, and all states s, t in $\mathcal{M}$:*

$$(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, t) \quad \text{iff} \quad (\mathcal{M}, s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M}, t).$$

2. *If $\mathbf{N}$ is a class of neighbourhood models in which the disjoint union of any two models is $\mathcal{L}_\Delta$-saturated, then for all $\mathcal{M}, \mathcal{M}'$ in $\mathbf{N}$,*

$$(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}', s') \quad \text{iff} \quad (\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s').$$

Proof *Item 1*: It is an immediate consequence of Lemma 5.4.15.

Item 2: $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}', s')$ implies $(\mathcal{M} + \mathcal{M}', s) \equiv_\Delta (\mathcal{M} + \mathcal{M}', s')$ since the inclusion morphisms are Δ-bisimulations. By item 1, $(\mathcal{M} + \mathcal{M}', s) \sim_\Delta^{\mathsf{betw}} (\mathcal{M} + \mathcal{M}', s')$, hence by definition, $(\mathcal{M}, s) \sim_\Delta (\mathcal{M}', s')$. □

As the disjoint union of two finite neighbourhood models is finite, and finite neighbourhood models are clearly $\mathcal{L}_\Delta$-saturated, we have an immediate corollary.

Corollary 5.4.17 *Over the class of finite neighbourhood models, $\mathcal{L}_\Delta$-equivalence implies Δ-bisimilarity.* ⊣

5.5 Frame class (un)definability

Modal logic formulas can be used to capture neighbourhood frame properties, for example, $\Box\varphi \wedge \Box\psi \to \Box(\varphi \wedge \psi)$ defines the class of frames in which the neighbourhood collections are closed under finite intersection [129]. In the section, we use Δ-bisimulations to demonstrate that $\mathcal{L}_\Delta$ is too weak to define some well-known frame classes. These results were already proved in [61, Prop. 7], but without the use of a bisimulation argument.

Definition 5.5.1 (Definability) *A frame class* $\mathbf{F}$ is $\mathcal{L}_\Delta$-definable *if there is a set $\Phi \subseteq \mathcal{L}_\Delta$ such that for all frames F, $F \in \mathbf{F}$ iff $F \models \Phi$.* ⊣

Let $\mathbf{M}$ be the class of (monotone) neighbourhood frames (S, ν) in which $\nu(s)$ is closed under supersets, for all $s \in S$. Let $\mathbf{C}$ be the class of neighbourhood frames (S, ν) in which $\nu(s)$ is closed under intersections, for all $s \in S$.

Example 5.5.2 *Consider the neighbourhood frames* $F_1 = (\{s_1, s_2\}, \nu_1)$ *and* $F_2 = (\{t_1, t_2\}, \nu_2)$ *where* $\nu_1(s_1) = \{\{s_2\}, \{s_1, s_2\}\}$, $\nu_2(t_1) = \{\emptyset, \{t_1\}\}$, *and* $\nu_1(s_2) = \nu_2(t_2) = \emptyset$. *These two frames are illustrated here:*

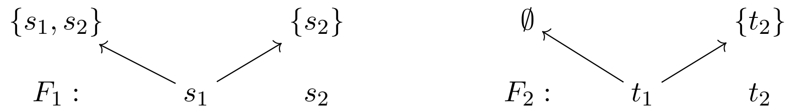

It can easily be checked that $Z = \{(s_1, t_1), (s_2, t_2)\}$ *is a* Δ*-bisimulation. The* Z*-coherent pairs are:* $(\emptyset, \emptyset)$, $(\{s_1\}, \{t_1\})$, $(\{s_2\}, \{t_2\})$ *and* $(\{s_1, s_2\}, \{t_1, t_2\})$. *The coherence condition for the pair* $(s_1, t_1) \in Z$ *and the* Z*-coherent pairs* $(\emptyset, \emptyset)$ *and* $(\{s_1, s_2\}, \{t_1, t_2\})$ *holds since* $\{s_1, s_2\} \in \nu_1(s_1)$ *and* $\emptyset \in \nu_2(t_1)$. *For* $(\{s_1\}, \{t_1\})$, *we have that* $\{s_1\}^c = \{s_2\} \in \nu_1(s_1)$ *and* $\{t_1\}^c = \{t_2\} \in \nu_2(t_1)$. *The case for* $(\{s_1\}, \{t_1\})$ *is clear. Since* $\nu_1(s_2) = \nu(t_2) = \emptyset$ *the coherence condition trivially holds for the pair* $(s_2, t_2) \in Z$. *Hence* Z *is a* Δ*-bisimulation. Note that* $F_1 \in \mathbf{M}$, *but* $F_2 \notin \mathbf{M}$. ⊣

Example 5.5.3 *Consider the neighbourhood frames* $F_3 = (\{s_1, s_2\}, \nu_3)$ *and* $F_4 = (\{t_1, t_2\}, \nu_4)$ *where* $\nu_3(s_1) = \{\{s_1, s_2\}, \{s_2\}, \emptyset\}$, $\nu_3(s_2) = \emptyset$, $\nu_4(t_1) = \{\{t_1\}, \{t_1, t_2\}, \{t_2\}\}$, *and* $\nu_4(t_2) = \emptyset$. *The two frames are illustrated here:*

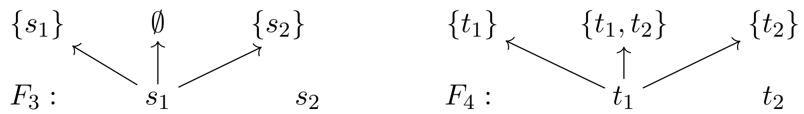

It can easily be checked that $Z = \{(s_1, t_1), (s_2, t_2)\}$ *is a* Δ*-bisimulation. Note that* $F_3 \in \mathbf{C}$, *but* $F_4 \notin \mathbf{C}$. ⊣

Proposition 5.5.4 *The frame classes* $\mathbf{M}$ *and* $\mathbf{C}$ *are not definable in* $\mathcal{L}_\Delta$. ⊣

Proof Example 5.5.2 shows that $\mathbf{M}$ is not $\mathcal{L}_\Delta$-definable, since suppose towards a contradiction that $\Phi \subseteq \mathcal{L}_\Delta$ defines $\mathbf{M}$. Then $F_1 \models \Phi$ and $F_2 \not\models \Phi$. Hence there is a valuation V_2 on F_2, a state t_j in F_2 and a $\varphi \in \Phi$ such that $(F_2, V_2), t_j \not\models \varphi$. We define a valuation V_1 on F_1 by $s_i \in V_1(p)$ iff $t_i \in V_2(p)$ for $i = 1, 2$ and all $p \in \mathsf{At}$. It follows that $((F_1, V_1), s_i) \sim_\Delta ((F_2, V_2), t_i)$ for $i = 1, 2$, and hence that $(F_1, V_1), s_j \not\models \varphi$, which implies that $F_1 \not\models \Phi$, a contradiction.

Similarly, Example 5.5.3 can be used to show that $\mathbf{C}$ is not $\mathcal{L}_\Delta$-definable. □

5.6 Characterisation Results

We first recall the basic definition of an ultrafilter. Let S be a nn-empty set. An *ultrafilter* over S is a collection of sets $\mathfrak{u} \subseteq \mathcal{P}(S)$ satisfying (i) $S \in \mathfrak{u}$ and $\emptyset \notin \mathfrak{u}$; (ii) $U_1, U_2 \in \mathfrak{u}$ implies $U_1 \cap U_2 \in \mathfrak{u}$; (iii) $U_1 \in \mathfrak{u}$ and $U_1 \subseteq U_2 \subseteq S$ implies $U_2 \in \mathfrak{u}$; and (iv) for all $U \subseteq S$ we have $U \in \mathfrak{u}$ or $U^c \in \mathfrak{u}$.

The collection of all ultrafilters over S will be denoted by $Ult(S)$. For $s \in S$, the principal ultrafilter generated by s is $\mathfrak{u}_s = \{U \subseteq S : s \in U\}$.

Definition 5.6.1 (Ultrafilter extension [89]) *Let* $\mathcal{M} = (S, \nu, V)$ *be a neighbourhood model. The* ultrafilter extension *of* $\mathcal{M}$ *is the triple* $\mathcal{M}^{ue} = (Ult(S), \nu^{ue}, V^{ue})$ *where* $V^{ue}(p) = \{\mathfrak{u} \in Ult(S) : V(p) \in \mathfrak{u}\}$ *and* $\nu^{ue} : Ult(S) \to \mathcal{P}(\mathcal{P}(Ult(S)))$ *is defined by*

$$\nu^{ue}(\mathfrak{u}) = \{\hat{U} \subseteq Ult(S) \ : \ U \subseteq S, \Box(U) \in \mathfrak{u}\}$$

where $\Box(U) = \{s \in S \ : \ U \in \nu(s)\}$ *and* $\hat{U} = \{\mathfrak{v} \in Ult(S) \mid U \in \mathfrak{v}\}$. ⊣

Lemma 5.6.2 *Let $(\mathcal{M}, s)$ be a pointed neighbourhood model. Then, $\mathcal{M}^{ue}$ is an $\mathcal{L}_\Delta$-saturated model and $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_s)$.* ⊣

Proof Since $\mathcal{L}_\Delta$ can be seen as a fragment of $\mathcal{L}_\Box$, [89, Lemma 4.24] ensures that $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_s)$ and [89, Proposition 4.25] ensures that $\mathcal{M}^{ue}$ is $\mathcal{L}_\Delta$-saturated. □

As in the $\mathcal{L}_\Box$ case, modal $\mathcal{L}_\Delta$-equivalence in a model implies Δ-bisimilarity in the ultrafilter extension.

Proposition 5.6.3 *Let $\mathcal{M}$ be a neighbourhood model and s, s' states in $\mathcal{M}$. Then,*

$$(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, s') \text{ implies } (\mathcal{M}^{ue}, \mathfrak{u}_s) \sim_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_{s'}).$$

Proof Suppose $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}, s')$. Lemma 5.6.2 ensures that $\mathcal{M}^{ue}$ is an $\mathcal{L}_\Delta$-saturated model and $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_s)$ and $(\mathcal{M}, s') \equiv_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_{s'})$. It follows that $(\mathcal{M}^{ue}, \mathfrak{u}_s) \equiv_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_{s'})$ and hence by Theorem 5.4.16(1), we have $(\mathcal{M}^{ue}, \mathfrak{u}_s) \sim_\Delta (\mathcal{M}^{ue}, \mathfrak{u}_{s'})$. □

We are now ready to prove the characterisation theorems.

Theorem 5.6.4 *An $\mathcal{L}_\Box$-formula is equivalent to an $\mathcal{L}_\Delta$-formula over the class of neighbourhood models iff it is invariant under Δ-bisimulation.* ⊣

Proof The left-to-right direction is clear. The proof of the converse is analogous to the proof of the characterisation result in [62, Theorem 4.4].

We will use the fact that classical modal logic **ML** is compact, i.e., for any set $\Phi \cup \{\varphi\}$ of $\mathcal{L}_\Box$-formulas, if $\Phi \models \varphi$, then there is a finite subset $\Phi_0 \subseteq \Phi$ such that $\Phi_0 \models \varphi$. The compactness of classical modal logic follows from strong completeness of the proof system $\mathbb{ML}$ of logic **ML** with respect to neighbourhood frames, cf. [37, Section 9.2]. Let $\varphi \in \mathcal{L}_\Delta$ be invariant under Δ-bisimulation over the class of all neighbourhood models, and let

$$\mathrm{MOC}(\varphi) = \{\psi^t \in \mathcal{L}_\Box \mid \varphi \models \psi,\, \psi \in \mathcal{L}_\Delta\}$$

be the translations of modal $\mathcal{L}_\Delta$-consequences of φ. By compactness of classical modal logic, it suffices to prove that $\mathrm{MOC}(\varphi) \models \varphi$. So let $(\mathcal{M}, s)$ be a pointed neighbourhood model such that $(\mathcal{M}, s) \models \mathrm{MOC}(\varphi)$. A standard compactness argument shows that $\mathrm{MOC}(\varphi) \cup \{\varphi\}$ is consistent, and hence satisfable in a pointed neighbourhood model $(\mathcal{M}', s')$, and we have $(\mathcal{M}, s) \equiv_\Delta (\mathcal{M}', s')$, since for all $\psi \in \mathcal{L}_\Delta$, $(\mathcal{M}, s) \models \psi$ implies that $\psi^t \in \mathrm{MOC}(\varphi)$ and hence $(\mathcal{M}', s') \models \psi$. Conversely, if $(\mathcal{M}, s) \not\models \psi$ then $(\neg\psi)^t \in \mathrm{MOC}(\varphi)$ and hence $(\mathcal{M}', s') \not\models \psi$.

Take now the disjoint union $\mathcal{M} + \mathcal{M}'$. It follows by Lemma 5.4.10 that $(\mathcal{M} + \mathcal{M}', s) \equiv_\Delta (\mathcal{M} + \mathcal{M}', s')$. Taking the ultrafilter extension of $\mathcal{M} + \mathcal{M}'$, Lemma 5.6.2 and Theorem 5.4.16 give us that $((\mathcal{M} + \mathcal{M}')^{ue}, \mathfrak{u}_s) \sim_\Delta ((\mathcal{M} + \mathcal{M}')^{ue}, \mathfrak{u}_{s'})$. We now have: From $(\mathcal{M}', s') \models \varphi$, it follows from Lemmas 5.4.10 and 5.6.2 that $((\mathcal{M} + \mathcal{M}')^{ue}, \mathfrak{u}_{s'}) \models \varphi$. As φ is invariant under Δ-bisimulations, $((\mathcal{M} + \mathcal{M}')^{ue}, \mathfrak{u}_s) \models \varphi$, and consequently $(\mathcal{M}, s) \models \varphi$, which concludes the proof that $\mathrm{MOC}(\varphi) \models \varphi$. □

In [89], a Van Benthem style characterisation theorem was given for classical modal logic with respect to a two-sorted first-order correspondence language $\mathcal{L}_1$. The two sorts **s** and **n** correspond to states and to neighbourhoods, respectively, and the basic idea of viewing a neighbourhood model as a first-order $\mathcal{L}_1$-structure is to encode the neighbourhood function ν as a relation

$R_\nu \subseteq \mathbf{s} \times \mathbf{n}$ between states and neighbourhoods, and encode subsets via the (inverse) element-of relation $R_\ni \subseteq \mathbf{n} \times \mathbf{s}$ between neighbourhoods and states. The language $\mathcal{L}_1$ is a first-order language with equality which contains a unary predicate symbol P (of sort $\mathbf{s}$) for each $p \in \mathsf{At}$, a binary relation symbol N (interpreted by R_ν), and a binary relation symbol E (interpreted by $R_\ni$). A translation $(-)^\sharp \colon \mathcal{L}_\Box \to \mathcal{L}_1$ is defined recursively over the Boolean connectives and atomic propositions, and by $(\Box\varphi)^\sharp = \exists u\,(\,x\mathsf{N}u \,\wedge\, \forall y\,(\,u\mathsf{E}y \leftrightarrow \varphi^\sharp\,))$. We refer to [89, section 5] for further details.

Theorem 5.6.5 *A first-order $\mathcal{L}_1$-formula is equivalent to an $\mathcal{L}_\Delta$-formula over the class of neighbourhood models iff it is invariant under Δ-bisimulation.* ⊣

Proof Let $\alpha \in \mathcal{L}_1$ be invariant under Δ-bisimulations. It follows from Lemma 5.4.10 that α is invariant under neighbourhood morphisms, and hence under behavioural equivalence. From the characterisation theorem [89, Theorem 5.5] it follows that α is equivalent to $\varphi^\sharp$ for some formula $\varphi \in \mathcal{L}_\Box$ which is necessarily also invariant under Δ-bisimulations. Hence by our Theorem 5.6.4, φ is equivalent to ψ^t for some $\psi \in \mathcal{L}_\Delta$. □

5.7 Craig Interpolation for Contingency Logic

In this section, we prove that contingency logic has the Craig interpolation property [43]. A logic has Craig interpolation property, if for all formulas φ, ψ such that $\models \varphi \to \psi$, there is an *interpolant* χ containing only common propositional variables of φ and ψ such that $\models \varphi \to \chi$ and $\models \chi \to \psi$. Craig interpolation is closely related to the expressive power of logics as it entails Beth's definability theorem [22]. Craig interpolation for modal logics has been explored from algebraic perspective in various studies [115, 116, 118, 74]. Recently, Seifan et al. in [147] studies interpolation from coalgebraic perspective. They propose a condition under which coalgebraic modal logic has uniform interpolation. It can be shown that contingency logic falls under their theorem, and hence it has uniform interpolation. Since uniform interpolation implies Craig interpolation, it follows that contingency logic has Craig interpolation. Here, we provide an elementary prove to show that classical contingency logic **CCL** enjoys Craig interpolation property using a similar approach as in [89], where the authors show a similar result for classical modal logic **ML** using ultrafilter extensions. We first need the following auxiliary definitions.

Let At be a set of atomic propositions and let $\mathsf{At}' \subseteq \mathsf{At}$. We denote by $\mathcal{L}_\Delta(\mathsf{At}')$ the sublanguage of $\mathcal{L}_\Delta$ generated by At'. Now, we generalise the notions we have so far to the sublanguage $\mathcal{L}_\Delta(\mathsf{At}')$. Let $\mathcal{M} = (S, \nu, V)$ and $\mathcal{M}' = (S', \nu', V')$ be neighbourhood models. This is completely straightforward, but for the sake of completeness, we provide the details. A binary relation $Z \subseteq S \times S'$ is a $\Delta^{\mathsf{At}'}$-bisimulation iff Z satisfies the **(Coherence)** condition (Def. 5.3.6), and for all $p \in \mathsf{At}'$, and $(s, s') \in Z$: $s \in V(p)$ iff $s' \in V'(p)$. A function $f : \mathcal{M} \to \mathcal{M}'$ is a $\Delta^{\mathsf{At}'}$-*morphism* between $\mathcal{M}$ and $\mathcal{M}'$ if $Gr(f)$ is a $\Delta^{\mathsf{At}'}$-bisimulation. A $\Delta^{\mathsf{At}'}$-congruence is defined as the kernel of a $\Delta^{\mathsf{At}'}$-morphism. We say two states $s \in \mathcal{M}$ and $s' \in \mathcal{M}'$ are $\mathcal{L}_\Delta(\mathsf{At}')$-equivalent, written $(\mathcal{M}, s) \equiv_{\mathcal{L}_\Delta(\mathsf{At}')} (\mathcal{M}', s')$, if s and s' satisfy the same $\mathcal{L}_\Delta(\mathsf{At})$-formulas. A subset $X \subseteq S$ is $\mathcal{L}_\Delta(\mathsf{At}')$-*compact* if for all sets Γ of $\mathcal{L}_\Delta(\mathsf{At}')$-formulas, if any finite subset $\Gamma' \subseteq \Gamma$ is satisfiable in X, then Γ is satisfiable in X. A model $\mathcal{M}$ is $\mathcal{L}_\Delta(\mathsf{At}')$-*saturated*, if for all $s \in S$ and all $\equiv_{\mathcal{L}_\Delta(\mathsf{At}')}$-closed neighbourhoods $X \in \nu(s)$, both X and X^c are $\mathcal{L}_\Delta(\mathsf{At}')$-compact.

Before presenting the main theorem, we first state a number of results that can be proved by retracing the arguments to At'.

Lemma 5.7.1 *Suppose* $\mathsf{At}' \subseteq \mathsf{At}$*, and let* $\mathcal{M} = (S, \nu, V)$*,* $\mathcal{M}' = (S', \nu', V')$ *be neighbourhood models. We have*

(1) Let $f : \mathcal{M} \to \mathcal{M}'$ be a $\Delta^{\mathsf{At}'}$-morphism. For all s in $\mathcal{M}$, and all $\varphi \in \mathcal{L}_\Delta(\mathsf{At}')$, we have $(\mathcal{M}, s) \models \varphi$ iff $(\mathcal{M}', f(s)) \models \varphi$.

*(2) Given an equivalence relation Z on S, Z is a $\Delta^{\mathsf{At}'}$-congruence on $\mathcal{M}$ iff Z satisfies the **(Coherence)** condition, and for all $(s,t) \in R$, and all $p \in \mathsf{At}'$: $s \in V(p)$ iff $t \in V(p)$.*

(3) If $\mathcal{M}$ is modally $\mathcal{L}_\Delta(\mathsf{At}')$-saturated, then all $\equiv_{\mathcal{L}_\Delta(\mathsf{At}')}$-closed subsets are definable by an $\mathcal{L}_\Delta(\mathsf{At}')$-formula.

(4) A binary relation Z on S is a $\Delta^{\mathsf{At}'}$-congruence iff Z is an equivalence relation and a nbh-$\Delta^{\mathsf{At}'}$-bisimulation.

(5) If $\mathcal{M}$ is modally $\mathcal{L}_\Delta(\mathsf{At}')$-saturated, then $\equiv_{\mathcal{L}_\Delta(\mathsf{At}')}$ is an $\Delta^{\mathsf{At}'}$-congruence.

(6) The ultrafilter extension $\mathcal{M}^{ue}$ of $\mathcal{M}$ is modally $\mathcal{L}_\Delta(\mathsf{At}')$-saturated. ⊣

Given a formula $\varphi \in \mathcal{L}_\Delta$, we let $\mathsf{At}(\varphi)$ denote the atomic propositions that occur in φ. Let $\varphi, \psi \in \mathcal{L}_\Delta$. Recall that for $\Gamma \cup \{\varphi\} \subseteq \mathcal{L}_\Delta$, we write $\Gamma \models \varphi$ if φ is a local semantic consequence of Γ. We write $\models \varphi$ if for all neighbourhood models $\mathcal{M}$ and all states s in $\mathcal{M}$, $(\mathcal{M}, s) \models \varphi$. Note that as every formula $\varphi \in \mathcal{L}_\Delta$ is equivalent to a formula $\varphi^t \in \mathcal{L}_\Box$, and classical modal logic **ML** is compact, it follows that contingency logic **CCL** is compact, as well. Now, we state the main result of this section.

Theorem 5.7.2 (Craig interpolation) *Let $\varphi_1, \varphi_2 \in \mathcal{L}_\Delta$. If $\models \varphi_1 \to \varphi_2$, then there exists a formula $\chi \in \mathcal{L}_\Delta$ with $\mathsf{At}(\xi) \subseteq \mathsf{At}(\varphi_1) \cap \mathsf{At}(\varphi_2)$ such that $\models \varphi_1 \to \chi$ and $\models \chi \to \varphi_2$.* ⊣

Proof We follow a similar method as in the proof of [89, Theorem 5.5.11]. Assume that $\models \varphi_1 \to \varphi_2$. For convenience of presentation, we let $\mathsf{At}_1 = \mathsf{At}(\varphi_1)$, $\mathsf{At}_2 = \mathsf{At}(\varphi_2)$, $\mathsf{At}_0 = \mathsf{At}_1 \cap \mathsf{At}_2$, and $\mathcal{L}_{\Delta i} = \mathcal{L}_\Delta(\mathsf{At}_i)$, $i = 0, 1, 2$. Denote by $\mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1) = \{\chi \in \mathcal{L}_{\Delta 0} \ : \ \varphi_1 \models \chi\}$ the set of $\mathcal{L}_{\Delta 0}$-consequences of φ_1. We claim that $\mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1) \models \varphi_2$. If we prove this claim, then by the compactness of contingency logic, there exist $\chi_1, \chi_2, ..., \chi_n \in \mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1)$ such that $\varphi_1 \models \chi_1 \wedge ... \wedge \chi_n$ and $\chi_1 \wedge ... \wedge \chi_n \models \varphi_2$, which implies that $\chi = \chi_1 \wedge ... \wedge \chi_n \in \mathcal{L}_{\Delta 0}$ is a Craig interpolant.

So, we will show that $\mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1) \models \varphi_2$. Let $(\mathcal{M}, s)$ be a pointed neighbourhood model such that $(\mathcal{M}, s) \models \mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1)$, and let $\Psi = \{\psi \in \mathcal{L}_{\Delta 0} \ : \ (\mathcal{M}, s) \models \psi\}$. We claim that $\Psi \cup \{\varphi_1\}$ is consistent. Assume, for the sake of contradiction, that it is inconsistent. Then, by compactness, for some finite subset $\{\psi_1, ..., \psi_n\} \subseteq \Psi$ we have $\models \psi_1 \wedge ... \wedge \psi_n \to \neg\varphi_1$. Hence, $\models \varphi_1 \to \neg\psi_1 \vee ... \vee \neg\psi_n$, but this implies $\neg\psi_1 \vee ... \vee \neg\psi_n \in \mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1)$, which contradicts the assumption that $(\mathcal{M}, s) \models \mathrm{Cons}_{\mathcal{L}_{\Delta 0}}(\varphi_1)$.

Hence, there is a pointed neighbourhood model $(\mathcal{M}', t)$ such that $(\mathcal{M}', t) \models \Psi \cup \{\varphi_1\}$. Due to the way we defined Ψ, $(\mathcal{M}, s) \equiv_{\mathcal{L}_{\Delta 0}} (\mathcal{M}', t)$. Now we take the disjoint union $\mathcal{M} + \mathcal{M}'$ with inclusion maps $\mathsf{in}_l : \mathcal{M} \to \mathcal{M} + \mathcal{M}'$ and $\mathsf{in}_r : \mathcal{M}' \to \mathcal{M} + \mathcal{M}'$. Since in_l and in_r preserve the truth of formulas, it follows that $(\mathcal{M} + \mathcal{M}', \mathsf{in}_l(s)) \equiv_{\mathcal{L}_{\Delta 0}} (\mathcal{M} + \mathcal{M}', \mathsf{in}_r(t))$. Now, we construct the ultrafilter extension of $\mathcal{M} + \mathcal{M}'$, which we denote by $\mathcal{U} = (U, \mu, V)$. By Lemma 5.6.2, we obtain that $(\mathcal{U}, \mathsf{u}_{\mathsf{in}_l(s)}) \equiv_{\mathcal{L}_{\Delta 0}} (\mathcal{U}, \mathsf{u}_{\mathsf{in}_r(t)})$. Since ultrafilter extensions are modally $\mathcal{L}_{\Delta 0}$-saturated (Lemma 5.7.1(6)), it follows from Lemma 5.7.1(5) that $\equiv_{\mathcal{L}_{\Delta 0}}$ is an Δ^{At_0}-congruence on $\mathcal{U}$. For notational convenience, let us in the rest of this proof denote by Z the relation $\equiv_{\mathcal{L}_{\Delta 0}}$ on $\mathcal{U}$. We observe that items 2 and 5 of Lemma 5.7.1 imply that Z is a Δ^{At_0}-bisimulation on $\mathcal{U}$. Now, define the neighbourhood function $\zeta : Z \to \mathcal{P}(\mathcal{P}(Z))$ for each $(u_l, u_r) \in Z$ by

$$\zeta\left((u_l, u_r)\right) = \{\pi_l^{-1}[C] \ : \ C \in \mu(u_l) \text{ or } C^c \in \mu(u_l)\} \cup \{\pi_r^{-1}[C] \ : \ C \in \mu(u_r) \text{ or } C^c \in \mu(u_r)\}$$

where $\pi_i : Z \to \mathcal{U}$ are the projection maps for $i = l, r$. We first show that π_l and π_r satisfy the **(Coherence)** condition, i.e.

$$\pi_i^{-1}[B] \in \zeta((u_l, u_r)) \text{ or } (\pi_i^{-1}[B])^c \in \zeta((u_l, u_r)) \iff B \in \mu(u_i) \text{ or } B^c \in \mu(u_i) \tag{5.8}$$

where $B \subseteq U$, and $i = l, r$. We only provide the proof for π_l, the proof for π_r is similar. To see this, let $(u_l, u_r) \in Z$, and let $(\pi_l^{-1}[B], B)$ be a $Gr(\pi_l)$-coherent pair (cf. Lemma 5.2.3(2)). The right to left direction of (5.8) is immediate. For the other direction, suppose that $\pi_l^{-1}[B] \in \zeta((u_l, u_r))$ or $(\pi_l^{-1}[B])^c \in \zeta((u_l, u_r))$. For now, let us assume $\pi_l^{-1}[B] \in \zeta((u_l, u_r))$. By the definition of ζ, the following cases can occur:

Case 1 $\pi_l^{-1}[B] = \pi_l^{-1}[C]$ for some $C \subseteq U$ with $C \in \mu(u_l)$ *or* $C^c \in \mu(u_l)$. If we prove that $B = C$, then we have $B \in \mu(u_1)$ *or* $B^c \in \mu(u_1)$, as desired. So, let us prove that $B = C$. Let $b \in B$. Since Z is an equivalence relation on $\mathcal{U}$, we have $(b, b) \in Z$. Therefore, $(b, b) \in \pi_l^{-1}[B]$, which implies that $(b, b) \in \pi_l^{-1}[C]$. We hence have that $b \in C$. Similarly, we can show that $C \subseteq B$.

Case 2 $\pi_l^{-1}[B] = \pi_r^{-1}[C]$ for some $C \subseteq U$ with $C \in \mu(u_r)$ *or* $C^c \in \mu(u_r)$. In this case, by Lemma 5.2.3(4) we have that (B, C) is Z-coherent, and since Z is a nbh-Δ^{At_0}-bisimulation (by Lemma 5.7.1(4)), it follows that $B \in \mu(u_l)$ *or* $B^c \in \mu(u_l)$.

Similarly, since $\pi_l^{-1}[B]^c = \pi_l^{-1}[B^c]$, we can show that (5.8) also holds for the case that $(\pi_l^{-1}[B])^c \in \zeta((u_l, u_r))$. We now define a valuation $\hat{V}$ on (Z, ζ) such that π_l and π_r become Δ^{At_1}-morphism and Δ^{At_2}-morphism, respectively. Let $p \in \mathsf{At}$, and $(u_l, u_r) \in Z$ we define

$$(u_l, u_r) \in \hat{V}(p) \iff \begin{cases} u_l \in V(p) & \text{if } p \in \mathsf{At}_1 \\ u_r \in V(p) & \text{if } p \in \mathsf{At}_2 \\ \text{never} & \text{if } p \in \mathsf{At} \setminus (\mathsf{At}_1 \cup \mathsf{At}_2) \end{cases}$$

The valuation $\hat{V}$ is well-defined due to Lemma 5.7.1(2). We now have: $(\mathcal{M}', t) \models \varphi_1$ implies that $(\mathcal{U}, \mathfrak{u}_{\mathsf{in}_r(t)}) \models \varphi_1$. As π_l is a Δ^{At_1}-morphism and preserves the truth of $\mathcal{L}_\Delta(\mathsf{At}_1)$-formulas, it follows that $(\mathcal{Z}, (\mathfrak{u}_{\mathsf{in}_l(t)}, \mathfrak{u}_{\mathsf{in}_r(s)})) \models \varphi_1$. So as (by assumption) $\models \varphi_1 \to \varphi_2$, we have that $(\mathcal{Z}, (\mathfrak{u}_{\mathsf{in}_l(t)}, \mathfrak{u}_{\mathsf{in}_r(s)})) \models \varphi_2$, and since π_r is an Δ^{At_2}-morphism, $(\mathcal{U}, \mathfrak{u}_{\mathsf{in}_l(s)}) \models \varphi_2$. This implies that $(\mathcal{M}, s) \models \varphi_2$, which completes the proof. □

5.8 Discussion and Future Work

We proposed a notion of contingency bisimulation on neighbourhood models, we related it to an existing notion of contingency bisimulation on Kripke models, and also provided a characterisation of (neighbourhood) contingency logic as a fragment of the modal logic of necessity, and of first-order logic. Our work contributes to a research program aiming at generalizing *knowing that* to *knowing whether, knowing how, knowing value*, etc. [156], including weaker modal notions than knowledge.

In [63], the logic **CL** was axiomatised over the class of all Kripke frames. We observed (below (5.4)) that the axiomatisation of **CL**, denoted by $\mathbb{CL}$, is sound and complete with respect to the class of augmented neighbourhood frames (which answers an open question in [62]). In [62] an axiomatisation $\mathbb{CCL}$ of classical contingency logic **CCL** is also given. This raises the questions of what the axiomatizations are of monotone contingency logic and regular contingency logic. Prop. 5.5.4 entails that one cannot fill these gaps with the axioms $\Delta\varphi \to \Delta(\varphi \to \psi) \vee \Delta(\neg\varphi \to \chi)$

and $\Delta(\psi \to \varphi) \wedge \Delta(\neg\psi \to \varphi) \to \Delta\varphi$ that are in $\mathbb{CL}$ but not in $\mathbb{CCL}$. So these questions remain open.

The **(Coherence)** condition in our definition of Δ-bisimulation is a non-local property, since one needs to check all Z-coherent pairs, so over large Kripke models the **Δ-Zig** and **Δ-Zag** conditions of o-Δ-bisimulations will be easier to check. As we proved that (over Kripke models) Δ-bisimilarity coincides with o-Δ-bisimilarity, one can view the **Δ-Zig** and **Δ-Zag** conditions as a back-forth characterisation of Δ-bisimilarity over Kripke models. We would like to find local zig-zag conditions also for Δ-bisimilarity over neighbourhood models.

The notion of Δ-bisimulation was based on the semantics of the modality Δ. It has a natural generalisation to the framework of coalgebraic modal logic [131, 40] in which the semantics of modalities is given by *predicate liftings*. In the next chapter 6, we work out this coalgebraic perspective. We will show that many of our results hold at this general coalgebraic level, and the notions of rel-Δ-bisimulation and nbh-Δ-bisimulation can be recovered by instantiating the generalised notion to Kripke and neighbourhood models for an appropriate choice of predicate liftings.

6

Bisimulation for weakly expressive coalgebraic modal logic

Contents

6.1 Introduction

In the last decades, modal logic [26] has found applications in different areas of computer science such as artificial intelligence, and verification of reactive and distributed systems. Different modal languages have been defined to fit specific semantic domains such as game frames [133], neighbourhood structures [121], probabilistic frames [56] and Markov chains [92]. Many definitions and results for these modal logics share some similar features, and it was therefore of theoretical interest to formulate a general framework in which modal logics and their semantic structures can be studied in a uniform setting. Such a framework is provided by coalgebra [141]. In coalgebra, state-based systems are defined parametric in the type of transitions and observations that the system can make. For example, a system could have nondeterministic or probabilistic transitions, and each state could have certain properties such as being accepting or not. The system type is formally specified by a functor. Given a functor T on the category of sets and functions, a *T-coalgebra* is a pair (X, γ) where X is a set and $\gamma \colon X \to TX$ is a function. By varying T, one obtains as concrete instances of T-coalgebras different systems such as automata, labelled transition systems, Markov chains, Krikpe models, neighbourhood models and much more [141, 82, 104].

Coalgebraic modal logic, as in [131, 103], is a framework in which modal logics for T-coalgebras can be developed parametric in the signature of the modal language and the coalgebra type functor T. Given a base logic (usually classical propositional logic), modalities are interpreted via so-called predicate liftings for the functor T. These are natural transformations that turn a predicate over the state space X into a predicate over TX. Given that T-coalgebras come with general notions of T-bisimilarity [141] and behavioural equivalence [104], coalgebraic modal logics are designed to respect those. In particular, if two states are behaviourally equivalent then they should satisfy the same formulas. If the converse holds, then the logic is said to be expressive. Such a result generalises the classic Hennessy-Milner theorem [93] which states that over the class of image-finite Kripke models, two states are Kripke bisimilar if and only if they satisfy the same formulas in Hennessy-Milner logic.

General conditions for when an expressive coalgebraic modal logic for T-coalgebras exists have been identified in [132, 23, 143]. A condition that ensures that a coalgebraic logic is expressive is when the set of predicate liftings chosen to interpret the modalities is *separating* [132]. Informally, a collection of predicate liftings is separating if they are able to distinguish non-identical elements from TX. This line of research in coalgebraic modal logic has thus taken as starting point the semantic equivalence notion of behavioral equivalence (or T-bisimilarity), and provided results for how to obtain an expressive logic. However, for some applications, modal logics that are not expressive are of independent interest. Such an example is given by contingency logic [62, 122] as we discussed in Chapter 5, where we proposed a notion of contingency bisimulation for contingency logic interpreted over neighbourhood models and proved the Hennessy-Milner theorem for it. We can now turn the question of expressiveness around and ask, given a modal language, whether we can generalise the notion of contingency bisimulation to the level of coalgebraic modal logic?

In this chapter, we propose a notion of Λ-bisimulation which is parametric in a collection Λ of predicate liftings, and therefore tailored to the expressiveness of a given coalgebraic modal logic. The definition relies on the notion of Z-coherent pairs, where Z is a relation between the state spaces of the relevant coalgebras. In particular, we see that if T is the neighbourhood functor and Λ consists of the usual neighbourhood modality, then Λ-bisimulation amounts to the notion of precocongruence for neighbourhood frames from [89]. We observe that coherent pairs have an abstract characterisation in terms of pullbacks and pushouts which makes it possible to prove most of our results using general category theoretical arguments. This suggests to us that Λ-bisimulations are a natural concept, which may be useful when considering coalgebraic modal logics over other categories than Sets. Moreover, we show that Λ-bisimulations, like T-bisimulations, form a complete lattice, and we show how they relate to T-bisimulations, behavioural equivalences and precocongruences. We also discuss their relationship to similar notions proposed by Gorin & Schröder [79] and Enqvist [54]. Our main result is a finitary Hennessy-Milner theorem (which does not assume Λ is separating): If T is finitary, then two states are Λ-bisimilar if and only if they satisfy the same modal Λ-formulas.

Overview. In Section 6.2 we fix notation and equip the reader with the necessary background material. In Section 6.3 we define our notion of Λ-bisimulation, study its properties, and relate it to other existing equivalence notions. Our Hennessy-Milner theorem is proved in Section 6.4.

The chapter concludes with a discussion of future and related work in Section 6.5.

6.2 Preliminaries

This section contains some basic definitions and results needed for reading this chapter and fixes notation.

6.2.1 Categories and functors

We briefly recall the basic concepts from category theory. An extensive introduction into category theory can be found in [107, 2].

A category C is a mathematical structure that consists of a collection $\mathbf{Obj}(\mathsf{C})$ of objects and a collection $\mathbf{Mor}(\mathsf{C})$ of morphisms (or arrows) between objects. Each morphism in $\mathbf{Mor}(\mathsf{C})$, written as $f : X \to Y$, has a *domain* object $X \in \mathbf{Obj}(\mathsf{C})$ and a *codomain* object $Y \in \mathbf{Obj}(\mathsf{C})$. In a category, morphisms with compatible codomain and domain can be composed associatively, and each object has an identity morphism. More precisely, given two morphisms $f : X \to Y$ and $g : Y \to Z$ in C, there is a morphism $g \circ f : X \to Z$ in C called the *composition* of f and g, and for each object X in C there is a morphism $\mathrm{Id}_X : X \to X$ such that for all $f : X \to Y$, $g : Y \to Z$ and $h : Z \to W$ in C, the following holds:

- $h \circ (g \circ f) = (h \circ g) \circ f$, and
- $f \circ \mathrm{Id}_X = \mathrm{Id}_Y \circ g$.

A functor $F : \mathsf{C} \to \mathsf{D}$ between categories C and D is a mapping that assigns to each object $X \in \mathbf{Obj}(\mathsf{C})$ an object $F(X) \in \mathbf{Obj}(\mathsf{D})$, and to each morphism $f : X \to Y \in \mathbf{Mor}(\mathsf{C})$ a morphism $F(f) : F(X) \to F(Y) \in \mathbf{Mor}(\mathsf{D})$ such that for all $X \in \mathbf{Obj}(\mathsf{C})$, $F(\mathrm{Id}_X) = \mathrm{Id}_{F(X)}$, and for all $f : X \to Y$ and $g : Y \to Z \in \mathbf{Mor}(\mathsf{C})$, $F(g \circ f) = F(g) \circ F(f)$.

In this chapter, we will work in the category Sets which has sets as objects and functions as morphisms. We will consider functors from Sets to Sets or to the opposite category $\mathsf{Sets}^{\mathrm{op}}$. The category $\mathsf{Sets}^{\mathrm{op}}$ is like Sets as it has sets as objects but an arrow f from X to Y in $\mathsf{Sets}^{\mathrm{op}}$ is a function $f : Y \to X$ in Sets. Note that a functor $F : \mathsf{Sets} \to \mathsf{Sets}^{\mathrm{op}}$ can also be seen as a functor $F^{op} : \mathsf{Sets}^{\mathrm{op}} \to \mathsf{Sets}$. We are mainly interested in the following three functors.

- The *covariant powerset* functor $\mathcal{P} : \mathsf{Sets} \to \mathsf{Sets}$ maps a set X to $\mathcal{P}(X)$, the powerset of X, and a function $f : X \to Y$ to the *direct image map*
$$\begin{aligned} \mathcal{P}(f) : \mathcal{P}(X) &\to \mathcal{P}(Y) \\ U &\mapsto f[U] = \{f(x) : x \in U\}. \end{aligned}$$
- The *contravariant powerset* functor $Q : \mathsf{Sets} \to \mathsf{Sets}^{\mathrm{op}}$ sends a set X to $\mathcal{P}(X)$ and a function $f : X \to Y$ to the inverse image map
$$\begin{aligned} Q(f) : Q(Y) &\to Q(X) \\ V &\mapsto f^{-1}[V] = \{x \in X : f(x) \in V\}. \end{aligned}$$
- The *neighbourhood* functor or *double contravariant* powerset functor $\mathcal{N} = Q^{\mathrm{op}}Q$ maps a set X to $\mathcal{N}(X) = Q^{\mathrm{op}}Q(X) : \mathsf{Sets} \to \mathsf{Sets}$ and a function $f : X \to Y$ to the function $\mathcal{N}(f) : \mathcal{N}(X) \to \mathcal{N}(Y)$ that is defined as follows:
$$\begin{aligned} \mathcal{N}(f) : \mathcal{N}(X) &\to \mathcal{N}(Y) \\ A &\mapsto \{B \subseteq Y : f^{-1}[B] \in A\}. \end{aligned}$$

We will see that in the theory of coalgebras, the following class of functors plays an important role.

Definition 6.2.1 (Finitary functor) *A functor* $T\colon \mathsf{Sets} \to \mathsf{Sets}$ *is said to be* finitary *if it satisfies for all sets* X

$$TX = \bigcup\{T\mathsf{in}_{X'}(TX') \subseteq TX : X' \subseteq X, X' \textit{ is finite}\}$$

where $\mathsf{in}_{X'}\colon X' \to X$ *is the inclusion map of* $X' \subseteq X$. ⊣

An example of finitary functor is the functor $\mathcal{P}_\omega\colon \mathsf{Sets} \to \mathsf{Sets}$ that maps a set X to $\mathcal{P}_\omega(X)$, the set of all its finite subsets, and a function $f\colon X \to Y$ to the function $\mathcal{P}_\omega f\colon \mathcal{P}_\omega X \to \mathcal{P}_\omega Y$ that is the restriction of $\mathcal{P}f$ to $\mathcal{P}_\omega X$. One can easily show that the powerset functor $\mathcal{P}$ and the neighbourhood functor $\mathcal{N}$ are not finitary.

Functors can themselves be viewed as objects in a category. The morphisms between them are called natural transformations. We will use natural transformations when defining coalgebraic semantics of modal logics. The definition is as follows. Given two categories C and D and two functors $F, G : \mathsf{C} \to \mathsf{D}$, a *natural transformation* $\upsilon\colon F \Rightarrow G$ is a family of morphisms in D

$$(\upsilon_X : FX \to GX)_{X \in \mathbf{Obj}(\mathsf{C})}$$

such that for all $f : X \to Y$ in C, $\upsilon_Y \circ F(f) = G(f) \circ \upsilon_X$, i.e., the following diagram commutes:

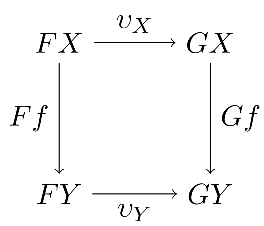

Figure 6.1. A natural transformation

6.2.2 Coalgebras

We assume that the reader is familiar with basic coalgebraic concepts. For a more thorough introduction to the theory of coalgebras, we refer to [141].

Definition 6.2.2 ($\boldsymbol{T}$-coalgebra) *Given a functor* $T\colon \mathsf{Sets} \to \mathsf{Sets}$*, a* T*-coalgebra is a pair* (X, γ) *with* $\gamma\colon X \to TX$*. The set* X *is called the* carrier *or the* state space*, and* γ *is called* the transition function*. A* T*-coalgebra morphism from* (X, γ) *to* (Y, δ)*, written* $f\colon (X, \gamma) \to (Y, \delta)$*, is a function* $f\colon X \to Y$ *such that* $Tf \circ \gamma = \delta \circ f$*, i.e., the following diagram commutes:*

The collection of T-coalgebras together with T-coalgebra morphisms form a category denoted by $Coalg(T)$. We will need some basic constructions in the category $Coalg(T)$ [141].

Definition 6.2.3 (Sub-coalgebra) *Given a* T*-coalgebra* $\mathbb{X} = (X, \gamma)$*, a* sub-coalgebra *of* $\mathbb{X}$ *is a* T*-coalgebra* $\mathbb{X}_0 = (X_0, \gamma_0)$ *such that* $X_0 \subseteq X$ *and the inclusion* $\mathsf{in}_{X_0}\colon X_0 \to X$ *is a* T*-coalgebra morphism from* $\mathbb{X}_0$ *to* $\mathbb{X}$. ⊣

The other construction that we will use further is the *coproduct* of T-coalgebras.

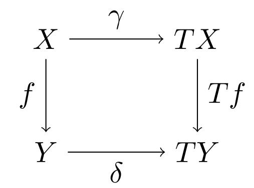

Figure 6.2. Coalgebra morphism

Definition 6.2.4 (Coproduct of T-coalgebras) *Given a functor $T\colon \mathsf{Sets} \to \mathsf{Sets}$, the coproduct of two T-coalgebras $\mathbb{X}_l = (X_l, \gamma_l)$ and $\mathbb{X}_r = (X_r, \gamma_r)$ is the T-coalgebra $\mathbb{X}_l + \mathbb{X}_r = (X_l + X_r, \zeta)$, where $X_l + X_r$ is the disjoint union of X_l and X_r, and $\zeta\colon X_l + X_r \to T(X_l + X_r)$ is defined by*

$$\zeta(x) = T\mathsf{in}_i \circ \gamma_i(x'). \qquad \text{(where } x = \mathsf{in}_i(x')\text{, for } i = l, r\text{)}$$

such that the inclusion maps $\mathsf{in}_l : X_i \to X_l + X_r$ and $\mathsf{in}_r\colon X_r \to X_l + X_r$ are T-coalgebra morphisms [141]. ⊣

6.2.3 Relations and coherence

We will use pullbacks and pushouts in what follows. We recall the general definitions and the concrete constructions in Sets.

Definition 6.2.5 (Pullbacks) *Let C be a category and let $f_l\colon X \to Z$ and $f_r\colon Y \to Z$ be morphisms in C. A* weak pullback *of f_l and f_r in C is a triple (B, g_l, g_r) where B is an object and $g_l\colon B \to X$, $g_r\colon B \to Y$ are morphisms in C such that $f_l \circ g_l = f_r \circ g_r$. Moreover, if B', $h_l\colon B' \to X$ and $h_r\colon B' \to Y$ are such that $f_l \circ h_l = f_r \circ h_r$, then there exists a morphism $m\colon B' \to B$ in C such that $h_l = g_l \circ m$ and $h_r = g_r \circ m$. The situation is depicted in Fig. 6.3(a). If the morphism m is unique, the triple (B, g_l, g_r) is called a* pullback. ⊣

In Sets, pullbacks can be obtained as follows. Given two functions $f_l\colon X \to Z$ and $f_r\colon Y \to Z$, we obtain a *pullback* of f_l and f_r by taking the triple $(B, g_l\colon B \to X, g_r\colon B \to Y)$, where

$$B = pb(f_l, f_r) = \{(x, y) \in X \times Y \mid f_l(x) = f_r(y)\}$$

and $g_l = \pi_l\colon B \to X$ and $g_r = \pi_r\colon B \to Y$ are the projections. The next definition recalls the dual notion of pullbacks.

Definition 6.2.6 (Pushouts) *Let C be a category and let $f_l\colon Z \to X$ and $f_r : Z \to Y$ be two morphisms in C. A* pushout *of f_l and f_r in C is a triple (P, p_l, p_r) where P is an object and $p_l\colon X \to P$ and $p_r\colon Y \to P$ are morphisms is C such that $p_l \circ f_l = p_r \circ f_r$. Moreover, if P', $k_l\colon X \to P'$ and $k_r\colon Y \to P'$ are such that $k_l \circ f_l = k_r \circ f_r$, then there exists a unique morphism $n : P \to P'$ in C such that $k_l = n \circ p_l$ and $k_r = n \circ p_r$, as illustrated in Figure 6.3(b).* ⊣

In Sets, given a relation $R \subseteq X \times Y$ the *pushout* of the projections $\pi_l\colon R \to X$ and $\pi_r\colon R \to Y$ is a triple $(P, p_l\colon X \to P, p_r\colon Y \to P)$ in which P is obtained concretely as follows. The relation R can be seen as a relation R_{X+Y} on the coproduct $X + Y$ by composing the projections with the coproduct injections $\mathsf{in}_l\colon X \to X + Y$ and $\mathsf{in}_r\colon Y \to X + Y$. More precisely, $R_{X+Y} = (\mathsf{in}_l \times \mathsf{in}_r)(R) = \{(\mathsf{in}_l(x), \mathsf{in}_r(y)) \mid (x, y) \in R\}$. Let $\overline{R}$ be the smallest equivalence relation on $X + Y$ that contains R_{X+Y}. Then we take $P = (X + Y)/\overline{R}$ to be the set of $\overline{R}$-equivalence classes with associated quotient map $q\colon X + Y \to P$, and we take $p_l = q \circ \mathsf{in}_l\colon X \to P$, $p_r = q \circ \mathsf{in}_r\colon Y \to P$.

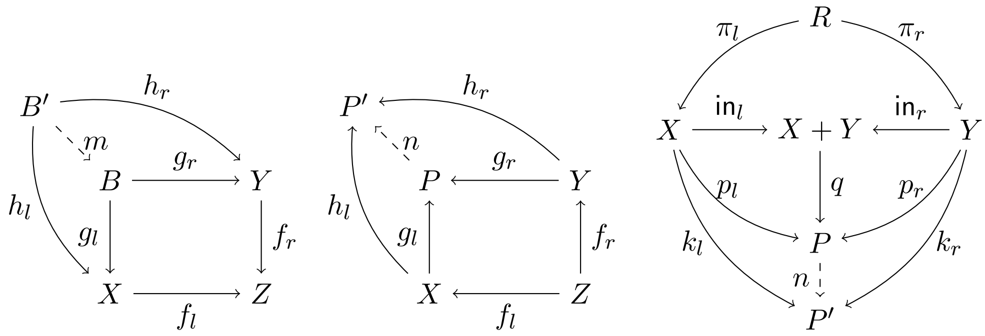

(a) Pullback (b) Pushout (c) Pushout in Sets

Figure 6.3. Pullback and pushout

An important property of functors in the theory of coalgebras is weak pullback preservation. A functor $T\colon \mathsf{Sets} \to \mathsf{Sets}$ preserves weak pullbacks if it maps weak pullbacks to weak pullbacks. More precisely, if whenever $(B, g_l\colon B \to X, g_r\colon B \to Y)$ is a weak pullback of $f_l\colon X \to Z$ and $f_r\colon Y \to Z$, then $(TB, Tg_l\colon TB \to TX, Tg_r\colon TB \to TY)$ is a weak pullback of $Tf_l\colon TX \to TZ$ and $Tf_r\colon TY \to TZ$. In terms of diagrams, it means that every weak pullback diagram on the left side is sent to a weak pullback on the right side:

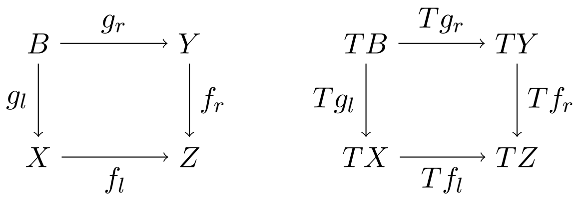

From the functors that we introduced so far, only the powerset functor preserves weak pullbacks.

Our definition of Λ-bisimulation relies on the notion of coherent pairs. For convenience, we recall the definition of coherent pairs from Chapter 5.

Definition 6.2.7 (*R*-coherent pairs) *Let $R \subseteq X \times Y$ be a relation with projections $\pi_l\colon R \to X$ and $\pi_r\colon R \to Y$, and let $U \subseteq X$ and $V \subseteq Y$. The pair (U, V) is R-*coherent *if $R[U] \subseteq V$ and $R^{-1}[V] \subseteq U$. In case $R \subseteq X \times X$ and $U \subseteq X$, then we say that U is R-*coherent *if (U, U) is R-coherent.* ⊣

We add an easy, but useful observation about coherent pairs in Lemma 6.2.8. Further properties of coherent pairs are found in Lemmas 5.2.3 and 5.2.4 (see page 106).

Lemma 6.2.8 *Let $R \subseteq X \times Y$ be a relation with projections $\pi_l\colon R \to X$ and $\pi_r\colon R \to Y$, and let $U \subseteq X$ and $V \subseteq Y$. The following are equivalent:*

1. *(U, V) is R-coherent.*
2. *(U, V) is in the pullback of $Q\pi_l$ and $Q\pi_r$.* ⊣

Due to Lemma 6.2.8(2), we will refer to $(pb(Q\pi_l, Q\pi_r), \pi'_l, \pi'_r)$ as the *pullback of R-coherent pairs.*

The following lemma shows that there is a fundamental connection between coherent pairs and pushouts of relations. It is also key in proving Propositions 6.3.11 and 6.3.12 later.

Lemma 6.2.9 *Let $R \subseteq X \times Y$ be a relation, and let (P, p_l, p_r) be the pushout of R. The triple (QP, Qp_l, Qp_r) is also a pullback of $(QR, Q\pi_l, Q\pi_r)$, and hence it is isomorphic to $(pb(Q\pi_l, Q\pi_r), \pi'_l, \pi'_r)$, the pullback of R-coherent pairs.* $\dashv$

Proof This lemma holds for the general reason that the contravariant powerset functor $Q\colon \mathsf{Sets} \to \mathsf{Sets}^{\mathrm{op}}$ is a left adjoint of itself, more precisely of $Q^{\mathrm{op}}\colon \mathsf{Sets}^{\mathrm{op}} \to \mathsf{Sets}$, and that left adjoints preserve colimits. Hence Q turns the pushout into a pullback. Since pullbacks are unique up to isomorphism, the result follows. The isomorphism is given concretely by the map $h\colon QP \to pb(Q\pi_l, Q\pi_r)$ defined for all $A \in Q(P)$ by $h(A) = (Qp_l(A), Qp_r(A))$. We verify that $(Q(p_l(A)), Q(p_r(A)))$ is R-coherent. So let $(x, y) \in R$. It follows that $p_l(x) = p_r(y)$, and hence $x \in Qp_l(A)$ iff $p_l(x) \in A$ iff $p_r(y) \in A$ iff $y \in Qp_r(A)$. To see that h is injective, suppose $A, A' \subseteq P$ and $a \in A \setminus A'$. The maps p_l and p_r are jointly surjective. If $a \in p_l[X]$, then there is a $x \in Qp_l(A)$ such that $p_l(x) = a$. If also $x \in Qp_l(A')$, then $p_l(x) = a \in A'$, a contradiction. Similarly, if $a \in p_r[Y]$, then there is a $y \in Qp_r(A)$ such that $p_r(y) = a$, and it must be the case that $y \notin Qp_r(A')$. Hence $h(A) \neq h(A')$. To see why h is surjective, it can be verified that if (U, V) is R-coherent, and we take $A \subseteq P$ to be $A = p_l[U] \cup p_r[V]$, then $h(A) = (U, V)$. For example, to see why $Qp_l(p_l[U]) = U$, first note that the inclusion $\supseteq$ always holds. Equality follows from the fact that (U, V) is R-coherent. Finally, we remark that (QP, Qp_l, Qp_r) is a competitor to the pullback of R-coherent pairs precisely because $(Qp_l(A), Qp_r(A))$ is R-coherent for all $A \subseteq P$. $\square$

6.2.4 Equivalence notions

We recall the equivalence notions in the theory of coalgebras, namely behavioural equivalence, (coalgebraic) bisimilarity, and precocongruence.

Definition 6.2.10 (Behavioural equivalence) *Let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ be T-coalgebras. Two states $x \in X$ and $y \in Y$ are* behaviourally equivalent *(notation: $\mathbb{X}, x \sim_{bh} \mathbb{Y}, y$), if there is a T-coalgebra $\mathbb{E} = (E, \epsilon)$ and a pair of T-coalgebra morphisms $f : \mathbb{X} \to \mathbb{E}$ and $g : \mathbb{Y} \to \mathbb{E}$ such that $f(x) = g(y)$. The situation is depicted below.*

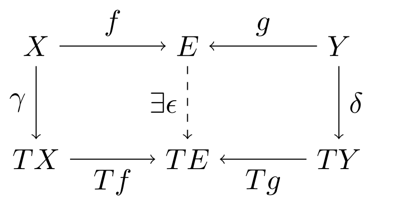

Figure 6.4. Behavioural equivalence

Definition 6.2.11 (T-bisimulation) *A relation $Z \subseteq X \times Y$ is a T-*bisimulation *between X and Y, if there exists a function $\zeta : Z \to TZ$ such that the projections $\pi_l : Z \to X$ and $\pi_r : Z \to Y$ are T-coalgebra morphisms, i.e., the following diagram commutes:*

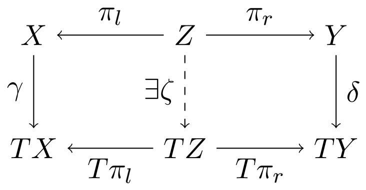

Figure 6.5. Z is a T-bisimulation.

Two states $x \in X$ and $y \in Y$ are T-bisimilar (notation: $\mathbb{X}, x \sim_T \mathbb{Y}, y$) if there is a T-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ linking x and y. ⊣

Definition 6.2.12 (Precocongruence) *Let $Z \subseteq X \times Y$ be a relation with pushout (P, p_l, p_r). The relation Z is a* precocongruence *between $\mathbb{X}$ and $\mathbb{Y}$ if there exists a function $\rho : P \to TP$ such that the pushout morphisms $p_l : X \to P$ and $p_r : Y \to P$ are T-coalgebra morphisms, i.e., if the following diagram commutes. If two states $x \in X$ and $y \in Y$ are related by some precocongruence, we write $\mathbb{X}, x \sim_p \mathbb{Y}, y$.*

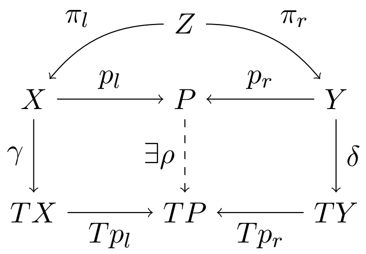

Figure 6.6. Z is a precocongruence.

We summarise a number of well-known facts on these notions. For any functor T, T-bisimilarity implies behavioural equivalence in T-coalgebras, however the converse only holds for the functors that preserve weak pullbacks [141, 1]. Every T-bisimulation is a precocongruence [89, Prop. 3.10(1)]. Moreover, on a single T-coalgebra, behavioural equivalence and precocongruences coincide [89, Theorem 3.12(2)].

6.2.5 Coalgebraic modal logic

Coalgebraic modal logic [131] is a uniform framework in which modal logics for coalgebras can be developed parametric in the type functor T and a choice of predicate lifting.

Syntax. A *similarity type* Λ is a set of modal operators with finite arities. We define the syntax of coalgebraic modal logic as follows.

Definition 6.2.13 *The set $\mathcal{L}_\Lambda$ of Λ-formulas is generated by the following grammar:*

$$\mathcal{L}_\Lambda \ni \varphi \quad ::= \quad \mathsf{t} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \heartsuit(\underbrace{\varphi, \ldots, \varphi}_{n\ times}) \quad (\heartsuit \in \Lambda, n\text{-}ary)$$

We use the standard definitions of the Boolean operators $\mathsf{f}, \vee$ and $\to$. ⊣

A T-coalgebraic semantics of $\mathcal{L}_\Lambda$-formulas is given by providing a Λ-*structure* $(T, (\llbracket\heartsuit\rrbracket)_{\heartsuit\in\Lambda})$ where T is a functor on Sets, and for each n-ary $\heartsuit \in \Lambda$, $\llbracket\heartsuit\rrbracket$ is an n-ary predicate lifting, i.e., $[\![\heartsuit]\!] : Q^n X \Rightarrow QT$ is a natural transformation, i.e., for every set X, it provides a function $[\![\heartsuit]\!]_X \colon Q^n X \to QTX$ such that for every mapping $f : X \to Y$ the following diagram commutes:

$$\begin{array}{ccc} Q^n X & \xrightarrow{[\![\heartsuit]\!]_X} & QTX \\ {\scriptstyle Qf}\uparrow & & \uparrow{\scriptstyle QTf} \\ Q^n Y & \xrightarrow[{[\![\heartsuit]\!]_Y}]{} & QTY \end{array}$$

Figure 6.7. Predicate lifting

where Q^n denotes the functor from $\mathsf{Sets} \times \ldots \mathsf{Sets}$ (n times) that maps a set X to $\mathcal{P}(X) \times \cdots \times \mathcal{P}(X)$ (n times), and maps every function $f \colon X \to Y$ to a function $Q^n f \colon Q^n Y \to Q^n Y$ which is defined by $Q^n f = Qf \times \ldots Qf$ (n times).

An n-ary modal operator $\heartsuit \in \Lambda$ is *monotone*, if it satisfies the following condition:

$$\text{For all } i = 1, \ldots, n \;\; A_i \subseteq B_i \subseteq X \quad \Rightarrow \quad [\![\heartsuit]\!]_X(A_1, \ldots, A_n) \subseteq [\![\heartsuit]\!]_X(B_1, \ldots, B_n).$$

Different choices of predicate liftings yield different Λ-structures and consequently different logics.

Semantics. Given a Λ-structure $(T, ([\![\heartsuit]\!])_{\heartsuit\in\Lambda})$, the truth of $\mathcal{L}_\Lambda$-formulas in a T-coalgebra $\mathbb{X} = (X, \gamma \colon X \to TX)$ is defined as follows:

$$\begin{array}{lll} \mathbb{X}, x \models \mathsf{t} & & \text{always} \\ \mathbb{X}, x \models \neg\varphi & \text{iff} & X, x \not\models \varphi \\ \mathbb{X}, x \models \varphi \wedge \psi & \text{iff} & \mathbb{X}, x \models \varphi \text{ and } \mathbb{X}, x \models \psi \\ \mathbb{X}, x \models \heartsuit(\varphi_1, \ldots, \varphi_n) & \text{iff} & \gamma(x) \in [\![\heartsuit]\!]_X([\![\varphi_1]\!]_{\mathbb{X}}, \ldots, [\![\varphi_n]\!]_{\mathbb{X}}). \end{array}$$

where $[\![\varphi]\!]_{\mathbb{X}} = \{x \in X \mid \mathbb{X}, x \models \varphi\}$ for all $\varphi \in \mathcal{L}_\Lambda$. Two states x in $\mathbb{X}$ and y in $\mathbb{Y}$ are (modally) $\mathcal{L}_\Lambda$-equivalent (notation: $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$), if they satisfy the same $\mathcal{L}_\Lambda$-formulas, i.e., $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$ if for all $\varphi \in \mathcal{L}_\Lambda$, $\mathbb{X}, x \models \varphi$ iff $\mathbb{Y}, y \models \varphi$.

A logical language $\mathcal{L}$ is *expressive* if (modally) $\mathcal{L}$-equivalent states are behaviourally equivalent [132]. Pattinson in [132] introduced the notion of a *separating* set of predicate liftings when studying expressive logics.

Definition 6.2.14 *A set $([\![\heartsuit]\!])_{\heartsuit\in\Lambda}$ of predicate liftings for a functor T is* separating *(for T) if every $t \in TX$ is uniquely determined by the set $\{((A_1, ..., A_n), \heartsuit) \in (\mathcal{P}(X))^n \times \Lambda \mid t \in [\![\heartsuit]\!]_X(A_1, ..., A_n)\}$. That is, if $t_1, t_2 \in TX$ and $t_1 \neq t_2$ then there is an n-ary $\heartsuit \in \Lambda$ and $A_1, \ldots, A_n \in \mathcal{P}(X)$ such that $t_1 \in [\![\heartsuit]\!]_X(A_1, \ldots, A_n)$ and $t_2 \notin [\![\heartsuit]\!]_X(A_1, \ldots, A_n)$, or vice versa.* $\dashv$

We provide some examples of modal languages and their coalgebraic semantics.

Example 6.2.15 Coalgebras for the covariant powerset functor $\mathcal{P}$ are Kripke frames. This is because every binary relation $R \subseteq X \times X$ can be presented as a function $R[-] \colon X \to \mathcal{P}(X)$ that maps every state $x \in X$ to $R[x]$, the set of R-successors of x. The similarity type $\Lambda = \{\Box\}$ for the basic modal language (without proposition letters) is given the usual Kripke semantics by interpreting $\Box$ via the predicate lifting

$$[\![\Box]\!]_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A\},$$

which is separating for $\mathcal{P}$, cf. [132].

Proposition letters can be included in the language by interpreting them as nullary predicate liftings. More precisely, given a set At of proposition letters, the basic modal language over At is obtained from the similarity type $\Lambda = \{\Box\} \cup \mathsf{At}$, This language is given its usual semantics in Kripke models which are coalgebras for the functor $TX = \mathcal{P}(X) \times \mathcal{P}(\mathsf{At})$ by taking the Λ-structure $(T, (\llbracket\heartsuit\rrbracket)_{\heartsuit\in\Lambda})$ where

$$\llbracket\Box\rrbracket_X(A) = \{(B,P) \in \mathcal{P}(X) \times \mathcal{P}(\mathsf{At}) \mid B \subseteq A\}, \text{ and}$$
$$\llbracket p\rrbracket_X(A) = \{(B,P) \in \mathcal{P}(X) \times \mathcal{P}(\mathsf{At}) \mid p \in P\}.$$

⊣

Example 6.2.16 The language of *contingency logic* (Def. 5.3.1) corresponds to the modal similarity type $\Lambda = \{\Delta\}$ and it is interpreted over Kripke frames (i.e. $\mathcal{P}$-coalgebras) via the predicate lifting

$$\llbracket\Delta\rrbracket_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A \textit{ or } B \subseteq A^c\}.$$

The predicate lifting $\llbracket\Delta\rrbracket$ is not separating for $\mathcal{P}$. To see this, consider the following example: let $X = \{x, y\}$, $B_1 = \{x\}$ and $B_2 = \{y\}$. For every subset $A \subseteq X$, $x \in A$ or $x \in A^c$ and $y \in A$ or $y \in A^c$. This means that for every subset $A \subseteq X$, $B_1 \subseteq A$ or $B_1 \subseteq A^c$ and $B_2 \subseteq A$ or $B_2 \subseteq A^c$. Hence, $\llbracket\Delta\rrbracket$ is not separating for $\mathcal{P}$. ⊣

Example 6.2.17 Neighbourhood frames are coalgebras for the functor $\mathcal{N}$. We obtain the neighbourhood semantics of the basic modal language, where $\Lambda = \{\Box\}$, by taking

$$\llbracket\Box\rrbracket_X(A) = \{B \in \mathcal{N}(X) \mid A \in B\},$$

which is separating for $\mathcal{N}$. ⊣

Example 6.2.18 The Neighbourhood semantics of contingency logic [61] is obtained by taking $T = \mathcal{N}$, $\Lambda = \{\Delta\}$, and

$$\llbracket\Delta\rrbracket_X(X) = \{B \in \mathcal{N}(X) \mid A \in B \textit{ or } A^c \in B\}.$$

As in the Kripke case, $\llbracket\Delta\rrbracket$ is not separating for $\mathcal{N}$. ⊣

Example 6.2.19 The language of *instantial neighbourhood logic (INL)* [19] arises from the similarity type $\Lambda = \{\Box_n \mid n \in \mathbb{N}\}$ where $\Box_n$ is $n+1$-ary for all $n \in \mathbb{N}$. The semantics of instantial neighbourhood logic is obtained by taking $T = \mathcal{P}\mathcal{P}$ and

$$\llbracket\Box_n\rrbracket_X(A_1, ..., A_n, B) = \{N \in \mathcal{P}(\mathcal{P}(X)) \mid \exists U \in N : U \subseteq B \text{ and for all } i = 1, \ldots, n : \ U \cap A_i \neq \emptyset\}.$$

The collection $\{\llbracket\Box_n\rrbracket \mid n \in \mathbb{N}\}$ is separating for $\mathcal{P}\mathcal{P}_\omega$: Suppose $N, N' \in \mathcal{P}(\mathcal{P}_\omega(X))$ and $B \in N \setminus N'$ with $B = \{x_1, \ldots, x_n\}$. Then $\llbracket\Box_n\rrbracket(\{x_1\}, \ldots, \{x_n\}, B)$ contains N, but not N'. It is not hard to see that any finite subset of $\{\llbracket\Box_n\rrbracket \mid n \in \mathbb{N}\}$ is not separating for $\mathcal{P}\mathcal{P}_\omega$. ⊣

6.3 Λ-bisimulation

In this section, we introduce the notion of Λ-bisimulation between T-coalgebras, and investigate its properties. This notion is parametric in the choice of a signature Λ and a Λ-structure $(T, (\llbracket\heartsuit\rrbracket)_{\heartsuit\in\Lambda})$. In the remaining part of the chapter, we therefore assume that we have fixed a functor $T\colon \mathsf{Sets} \to \mathsf{Sets}$, and for each $\heartsuit \in \Lambda$, a predicate lifting $\llbracket\heartsuit\rrbracket$ of appropriate arity. From now on, by abuse of language, we will also refer to Λ as the set of these predicate liftings. Moreover, we let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ denote arbitrary T-coalgebras.

6.3.1 Definition and basic properties

The definition of Λ-bisimulation is as follows.

Definition 6.3.1 (Λ-bisimulation) *Let $Z \subseteq X \times Y$ be a relation and let $(pb(Q\pi_l, Q\pi_r), \overline{\pi}_l, \overline{\pi}_r)$ be the associated pullback of Z-coherent pairs. The relation Z is a* Λ-bisimulation *between $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$, if for all $\heartsuit \in \Lambda$, with $\heartsuit$ n-ary, the following diagram commutes:*

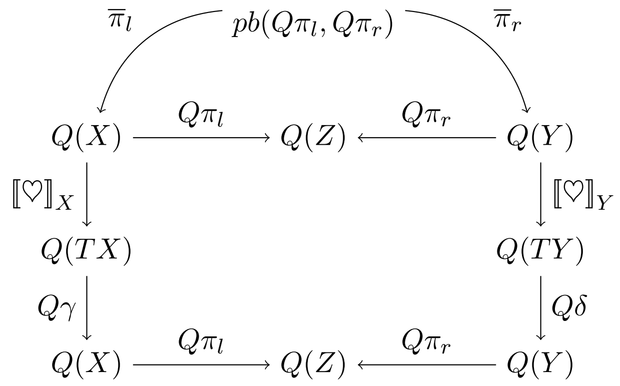

Figure 6.8. Λ-bisimulation

I.e., the following equality holds:

$$Q\pi_l \circ Q\gamma \circ [\![\heartsuit]\!]_X \circ \overline{\pi}_l^n = Q\pi_r \circ Q\delta \circ [\![\heartsuit]\!]_Y \circ \overline{\pi}_r^n \tag{6.1}$$

where $\overline{\pi}_l^n \colon pb(Q\pi_l, Q\pi_r)^n \to (QX)^n$ and $\overline{\pi}_r^n \colon pb(Q\pi_l, Q\pi_r)^n \to (QY)^n$ are the pointwise projections, for example, $\overline{\pi}_l((U_1, V_1), \ldots, (U_n, V_n)) = (U_1, \ldots, U_n)$.

In other words, the relation Z is a Λ-bisimulation if whenever $(x, y) \in Z$, then for all $\heartsuit \in \Lambda$, n-ary, and all Z-coherent pairs $(U_1, V_1), \ldots, (U_n, V_n)$, we have that

$$\gamma(x) \in [\![\heartsuit]\!]_X(U_1, \ldots, U_n) \quad \textit{iff} \quad \delta(y) \in [\![\heartsuit]\!]_Y(V_1, \ldots, V_n). \tag{\textbf{Coherence}}$$

We write $\mathbb{X}, x \sim_\Lambda \mathbb{Y}, y$, when there is a Λ-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ that contains (x, y). A Λ-bisimulation on a T-coalgebra $\mathbb{X}$ is a Λ-bisimulation between $\mathbb{X}$ and $\mathbb{X}$. We say two states x in $\mathbb{X}$ and y in $\mathbb{Y}$ are Λ-bisimilar *(notation: $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{Y}, y$) if $\mathbb{X} + \mathbb{Y}, \mathsf{in}_l(x) \sim_\Lambda \mathbb{X} + \mathbb{Y}, \mathsf{in}_r(y)$.* ⊣

Remark 6.3.2 *We note that using the basic observation (Lemma 6.2.9) that the coherent pairs is isomorphic to the dual of the pushout, Λ-bisimulations can be reformulated in terms of complex algebras (i.e. algebras for the modal signature) as follows: A relation Z between two T-coalgebras $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ is a Λ-bisimulation iff the pullback of Z-coherent pairs $(pb(Q\pi_l, Q\pi_r), Q\pi_l, Q\pi_r)$ is a congruence between the complex algebras $\mathbb{X}^* = (LQX, \gamma^*)$ and $\mathbb{Y}^* = (LQY, \delta^*)$, where $L : \mathsf{BA} \to \mathsf{BA}$ is the functor corresponding to the modal signature Λ, i.e., $LA = \coprod_{\heartsuit \in \Lambda} A^{ar(\heartsuit)}$, $\gamma^* = LQX \xrightarrow{\sigma_X} QTX \xrightarrow{Q\gamma} QX$, $\delta^* = LQY \xrightarrow{\sigma_Y} QTY \xrightarrow{Q\delta} QY$, and $\sigma \colon LQ \Longrightarrow QT$ is the bundling up of all predicate liftings into one natural transformation.* [4] ⊣

We define Λ-bisimilarity via the coproduct for the same reason that rel-Δ-bisimilarity was defined via the coproduct in Chapter 5 (Def. 5.3.6). Namely, rel-Δ-bisimulations are instances of Λ-bisimulations, as we will see later in Example 6.2.16, and in Example 5.3.7(4) we saw that even

[4] This observation has been made by Alexander Kurz during the Ph.D. defense of the author of the thesis.

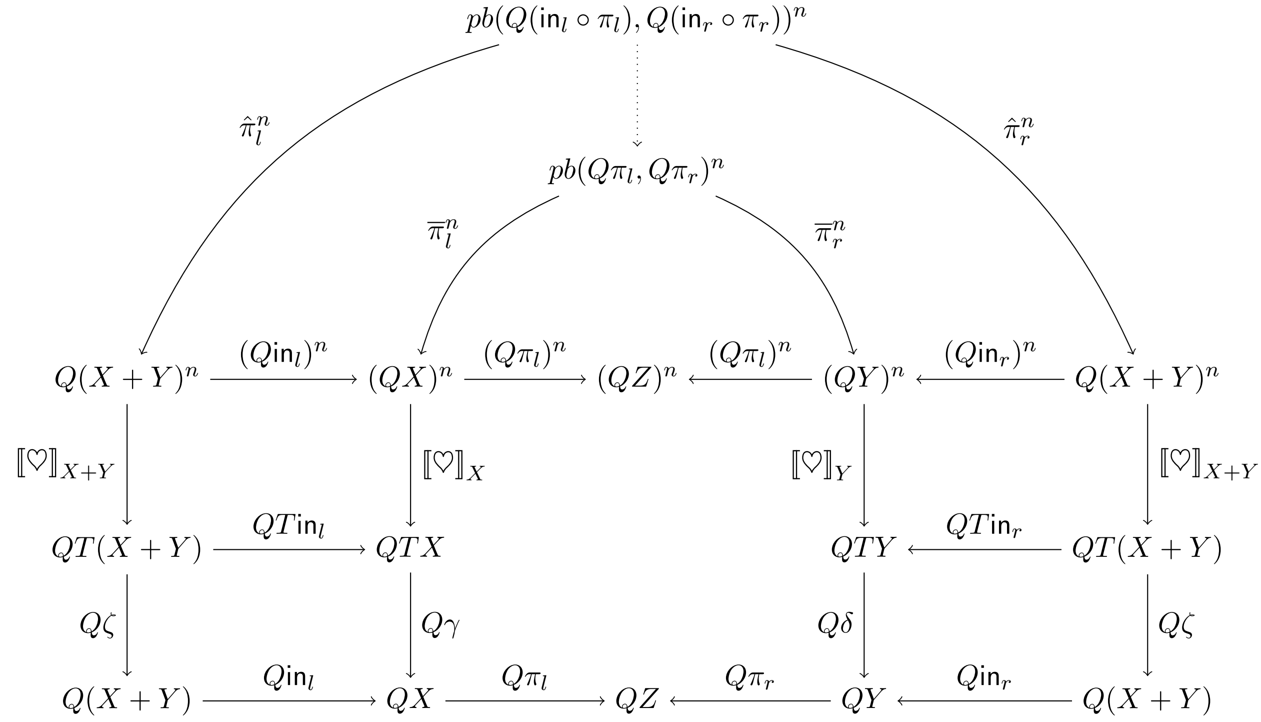

Figure 6.9. Diagram of the proof of Prop. 6.3.3

between $\mathcal{P}_\omega$-coalgebras, it is possible for two states to satisfy the same modal $\mathcal{L}_\Delta$-formulas without being linked by a rel-Δ-bisimulation. So in order to obtain a Hennessy-Milner theorem over finite models with respect to Λ-bisimilarity, we cannot define Λ-bisimilarity using the $\sim_\Lambda$ notion.

In the next proposition we will show that over a single T-coalgebra, $\sim_\Lambda$ and $\sim_{\Lambda_+}$ coincide, however in general they differ.

Proposition 6.3.3 *For all $x, x' \in X$ and $y \in Y$,*

1. $\mathbb{X}, x \sim_\Lambda \mathbb{Y}, y$ *implies* $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{Y}, y$. *The implication is strict.*
2. $\mathbb{X}, x \sim_\Lambda \mathbb{X}, x'$ *iff* $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{X}, x'$. ⊣

Proof *Item 1.* Let $Z \subseteq X \times Y$ be a Λ-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$. We show that the relation $(\mathsf{in}_l \times \mathsf{in}_r)(Z) = \{(\mathsf{in}_l(x), \mathsf{in}_r(y)) \mid (x,y) \in Z\}$ is a Λ-bisimulation on $\mathbb{X} + \mathbb{Y} = (X+Y, \zeta)$. The proof follows from the commutativity of the diagram in Figure 6.9 where $\heartsuit \in \Lambda$ is arbitrary. The commutativity follows from observing that $pb(Q(\mathsf{in}_l \circ \pi_l), Q(\mathsf{in}_r \circ \pi_r))$ with $\hat{\pi}_l \circ Q\mathsf{in}_l$ and $\hat{\pi}_r \circ Q\mathsf{in}_r$ is a competitor to the pullback $(pb(Q\pi_l, Q\pi_r), \overline{\pi}_l, \overline{\pi}_r)$. This yields a mediating map (dashed arrow) such that the upper part of the diagram commutes. The lower, outer parts commute due to naturality of $[\![\heartsuit]\!]$ and the inclusions being T-coalgebra morphisms.

Item 2. The direction from left to right follows from item 1. For the other direction, assume that Z is a Λ-bisimulation on $\mathbb{X} + \mathbb{X} = (X+X, \zeta)$. We show that $Z' = \{(w, w') \in X \times X \mid \exists i, j \in \{r, l\} : (\mathsf{in}_i(w), \mathsf{in}_j(w')) \in Z\}$ is a Λ-bisimulation on $\mathbb{X}$. First, note that is (U, V) if Z'-coherent, then $(U+U, V+V)$ is Z-coherent. Let $(x, x') \in Z'$, then $(\mathsf{in}_i(x), \mathsf{in}_j(x')) \in Z$, for some $i, j \in \{l, r\}$. Since Z is a Λ-bisimulation, it follows that:

$$\zeta(\mathsf{in}_i(x)) \in [\![\heartsuit]\!]_{X+X}(U+U) \iff \zeta(\mathsf{in}_j(x')) \in [\![\heartsuit]\!]_{X+X}(V+V) \qquad (6.2)$$

To complete the proof, it remains to show that for very $U \subseteq X$

$$\gamma(x) \in [\![\heartsuit]\!]_X[U] \iff \zeta(\mathsf{in}_i(x)) \in [\![\heartsuit]\!]_{X+X}(U+U) \ \ (i = l, r) \tag{6.3}$$

But this follows from naturality of $[\![\heartsuit]\!]$ and the fact that inclusion maps are T-coalgebra morphism. Item 2 then follows from (6.2) and (6.3). $\square$

The next lemma provides an easy observation about dual modal operators that we will use further in the examples.

Lemma 6.3.4 *Let $\heartsuit, \heartsuit' \in \Lambda$ be two n-ary dual modalities, that is $\heartsuit = \neg\heartsuit'\neg$. A relation Z is a $\heartsuit$-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ iff Z is a $\heartsuit'$-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$.* $\dashv$

Proof First, $\heartsuit = \neg\heartsuit'\neg$ means that for all sets W, and all $A_1, \dots, A_n \subseteq W$, we have that $[\![\heartsuit]\!]_W(A_1, \dots, A_n) = ([\![\heartsuit']\!]_W(A_1^c, \dots, A_n^c))^c$. We note that if $Z \subseteq X \times Y$, $U \subseteq X$ and $V \subseteq Y$, then the pair (U, V) is Z-coherent iff (U^c, V^c) is Z-coherent. Hence, $\gamma(x) \in [\![\heartsuit]\!]_X(U_1, ..., U_n)$ iff $\gamma(x) \notin [\![\heartsuit']\!]_X(U_1^c, ..., U_n^c)$ iff $\delta(y) \notin [\![\heartsuit']\!]_Y(V_1^c, ..., V_n^c)$ iff $\delta(y) \in [\![\heartsuit]\!]_Y(V_1, ..., V_n)$. $\square$

We provide some examples of our notion of Λ-bisimulation.

Example 6.3.5 Taking $T = \mathcal{P}$ (i.e. T-coalgebras are Kripke frames) and $\Lambda = \{\Box\}$ (or $\Lambda = \{\Diamond\}$), then a relation Z between Kripke frames $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ is a Λ-bisimulation if for all $(x, y) \in Z$ and all Z-coherent pairs (U, V): $\gamma(x) \subseteq U$ iff $\delta(y) \subseteq V$. An easy proof shows that if Z is a Kripke bisimulation then Z is a Λ-bisimulation. However, a Λ-bisimulation may not be a Kripke bisimulation. Consider the following Kripke frames: $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$, where $X = \{x, x_1, x_2\}$, $\gamma(x) = \{x_1\}$, $Y = \{y, y_1, y_2\}$ and $\delta(y) = \{y_1, y_2\}$. The relation $Z = \{(x, y), (x_1, y_1), (x_2, y_1), (x_2, y_2)\}$ is a Λ-bisimulation, but it is not a Kripke bisimulation, since the successor y_2 of y is not related to a successor of x. The situation is depicted below, where Z is indicated by dashed lines.

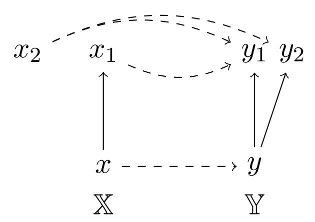

Still, when considering the associated bisimilarity notions, we find that Λ-bisimilarity coincides with Kripke bisimilarity. This follows from our Proposition 6.3.13, later, using that $\Box$ (and $\Diamond$) is separating and $\mathcal{P}$ preserves weak pullbacks.

This choice of T and Λ demonstrates that, in general, Λ-bisimulations are not closed under relational composition. To see this, let $\mathbb{X} = (X, \gamma)$, $\mathbb{Y} = (Y, \delta)$ and $\mathbb{W} = (W, \alpha)$ be the three Kripke frames depicted below together with the two relations $Z_1 \subseteq X \times Y$ and $Z_2 \subseteq Y \times W$ (indicated by dashed lines): It is straightforward to check that Z_1 and Z_2 are Λ-bisimulations, but the composition $Z_1; Z_2 = \{(x, w)\}$ is not, because $(\{x, x_1\}, \{w\})$ is $Z_1; Z_2$-coherent and $\gamma(x) \nsubseteq \{x, x_1\}$ and $\gamma(x) \nsubseteq \{x_2\}$, whereas $\alpha(w) \subseteq \{w\}$.

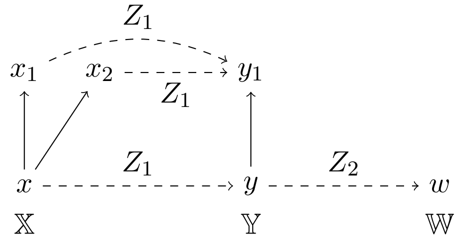

⊣

Example 6.3.6 Taking $T = \mathcal{N}$ (i.e. neighbourhood frames) and $\Lambda = \{\Box\}$, where $\Box$ is the neighbourhood modality from Example 6.2.17, we find that a relation Z is a Λ-bisimulation between neighbourhood frames $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ if for all $(x, y) \in Z$ and all Z-coherent (U, V): $U \in \gamma(x)$ iff $V \in \delta(y)$. This shows that Λ-bisimulations are the same as *precocongruences* (Def. 6.2.12) due to [89, Proposition 3.16]. We will discuss the relation between precocongruences and Λ-bisimulations further in subsection 6.3.2. ⊣

Example 6.3.7 Taking $T = \mathcal{P}$ and $\Lambda = \{\Delta\}$, where Δ is the contingency modality from Example 6.2.16, then Z is a Λ-bisimulation between Kripke frames $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ if for all $(x, y) \in Z$ and all Z-coherent (U, V): $\gamma(x) \subseteq U$ *or* $\gamma(x) \subseteq U^c$ iff $\delta(y) \subseteq V$ *or* $\delta(y) \subseteq V^c$. This is exactly the definition of a *rel-Δ-bisimulation* which was introduced in Def. 5.3.6 on page 109. ⊣

Example 6.3.8 Taking $T = \mathcal{N}$ and $\Lambda = \{\Delta\}$, where Δ is the neighbourhood contingency modality from Example 6.2.18, then by instantiating (**Coherence**) for Δ, we have that Z is a Λ-bisimulation between neighbourhood frames $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ if for all $(x, y) \in Z$ and all Z-coherent (U, V): $U \in \gamma(x)$ *or* $U^c \in \gamma(x)$ iff $V \in \delta(y)$ *or* $V^c \in \delta(y)$. This is exactly the definition of a *nbh-Δ-bisimulation* which was introduced in Def. 5.4.4 on page 114. ⊣

The following proposition shows that Λ-bisimulations enjoy many of the properties known to hold for Kripke bisimulations. In particular, even though Λ-bisimulations do not need to be closed under composition (cf. Example 6.3.5), we can still show that on a single T-coalgebra, $\sim_\Lambda$ is an equivalence relation.

For the sake of a smooth presentation, in the reminder of this chapter we assume that the predicate liftings are unary, unless we specify otherwise.

Proposition 6.3.9 *Let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ be T-coalgebras.*

1. *The identity relation* Id $\subseteq X \times X$ *is a Λ-bisimulation on* $\mathbb{X}$.
2. *If $Z \subseteq X \times Y$ is a Λ-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$ then $Z^{-1} \subseteq Y \times X$ is a Λ-bisimulation between $\mathbb{Y}$ and $\mathbb{X}$.*
3. *The set of Λ-bisimulations is closed under arbitrary unions: If $Z_i \subseteq X \times Y$, $i \in I$, are Λ-bisimulations, then so is $\bigcup_{i \in I} Z_i$.*
4. *The relation $\sim_\Lambda$ is the largest Λ-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$.*
5. *The relation $\sim_\Lambda$ on $\mathbb{X}$ is an equivalence relation.* ⊣

Proof Item 1 is straightforward to check. We omit the details.
Item 2 follows from the fact that for every relation $R \subseteq X \times Y$, the pair (U, V) is R-coherent if and only if (V, U) is R^{-1}-coherent.
Item 3 : Let $Z_i \subseteq X \times Y$, $i \in I$, be Λ-bisimulations, and let $Z = \bigcup_{i \in I} Z_i$. To show that Z is a Λ-bisimulation, assume that $(x, y) \in Z$, $\heartsuit \in \Lambda$, and (U, V) is a Z-coherent pair. From $(x, y) \in Z$ it follows that $(x, y) \in Z_i$ for some $i \in I$, and since $Z_i \subseteq Z$ we also have that (U, V) is Z_i-coherent. Hence

$$\gamma(x) \in [\![\heartsuit]\!]_X(U) \iff \delta(y) \in [\![\heartsuit]\!]_X(V).$$

This implies that Z is a Λ-bisimulation.
Item 4 is an immediate consequence of item 3.
Item 5 : We show that if Z is a Λ-bisimulation on $\mathbb{X}$, then the equivalence closure of Z is again a Λ-bisimulation on $\mathbb{X}$, which suffices due to item 4. So let Z be a Λ-bisimulation on $\mathbb{X}$. By items 1 and 2, we may assume that Z is reflexive and symmetric. The result follows by showing that the transitive closure $Z^+ = \bigcup_{n \geq 1} Z^n$ is a Λ-bisimulation. Due to item 3 it suffices to show that for all $n \geq 1$, Z^n is a Δ-bisimulation. The proof is by induction on n. The base case ($n = 1$) holds by assumption on Z. Assume it holds for n. To prove the inductive step, we first note that if (U, U'') is Z^{n+1}-coherent, then since Z^{n+1} is reflexive, it follows that $U = U''$. Now suppose $(x, x') \in Z^n$, $(x', x'') \in Z$ and (U, U) is Z^{n+1}-coherent. Since Z and Z^n are reflexive and $Z^{n+1} = Z^n; Z$, it follows that $Z \subseteq Z^{n+1}$ and $Z^n \subseteq Z^{n+1}$, and hence (U, U) is Z-coherent as well as Z^n-coherent. We then have

$$\begin{aligned} \gamma(x) \in [\![\heartsuit]\!]_X(U) &\iff \gamma(x') \in [\![\heartsuit]\!]_X(U) \quad \text{(by induction hypothesis)} \\ &\iff \gamma(x'') \in [\![\heartsuit]\!]_X(U) \quad \text{(since } Z \text{ is a } \Lambda\text{-bisimulation).} \end{aligned}$$

Hence Z^{n+1} is a Λ-bisimulation which concludes the proof. □

Λ-bisimulations were designed to match the expressiveness of the modal language. In the next proposition we show that indeed, Λ-bisimilar states satisfy the same $\mathcal{L}_\Lambda$-formulas.

Proposition 6.3.10 *If $\mathbb{X}, x \sim_\Lambda \mathbb{Y}, y$ then $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$.* ⊣

Proof Let $\mathbb{X}, x \sim_\Lambda \mathbb{Y}, y$, so there exists a Λ-bisimulation $Z \subseteq X \times Y$ such that $(x, y) \in Z$. The proof is by induction on φ. The only interesting part is the modal case of the inductive step. Assume that φ is of the form $\heartsuit\psi$. By induction hypothesis, $([\![\psi]\!]_\mathbb{X}, [\![\psi]\!]_\mathbb{Y})$ is Z-coherent. Since Z is a Λ-bisimulation, we have $\gamma(x) \in [\![\heartsuit]\!]_X([\![\psi]\!]_\mathbb{X})$ iff $\delta(y) \in [\![\heartsuit]\!]_Y([\![\psi]\!]_\mathbb{Y})$, which means that $\mathbb{X}, x \models \heartsuit\psi$ iff $\mathbb{Y}, y \models \heartsuit\psi$. □

6.3.2 Comparison with other notions

In this part, we compare our notion of Λ-bisimulation to the equivalence notions that are defined in Subsection 6.2.4, namely behavioural equivalence (Def. 6.2.10), T-bisimulations (Def. 6.2.11) and precogongruences (Def. 6.2.12). It turns out that Λ-bisimulations are closest to precocongruences. Finally, we also compare our notion to another similar proposal by Gorín and Schröder [79].

In the following proposition we give the first comparison between precocongruences, T-bisimulations and Λ-bisimulations.

Proposition 6.3.11 *Let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ be T-coalgebras, and Z be a relation between X and Y.*

1. *If Z is a T-bisimulation then Z is a Λ-bisimulation.*

2. *If Z is a precocongruence then Z is a Λ-bisimulation.* $\dashv$

Proof Item 1. Apply Q to the diagram of T-bisimulation (Figure 6.5), and take the pullback of $Q\pi_l$ and $Q\pi_r$. Then, by naturality of $[\![\heartsuit]\!]$, and the fact that π_l and π_r are coalgebra morphisms, the diagram in Figure 6.10(a) commutes and hence, Z is a Λ-bisimulation.
Item 2. Let $Z \subseteq X \times Y$ be a precocongruence relation with pushout (P, p_l, p_r), and let (U, V) be Z-coherent. By Lemma 6.2.9, there is a map $g\colon pb(Q\pi_l, Q\pi_r) \to QP$ such that $Qp_l \circ g = \overline{\pi}_l$ and $Qp_r \circ g = \overline{\pi}_r$. Then, by naturality of $[\![\heartsuit]\!]$ and the fact that p_l and p_r are T-coalgebra morphisms, it follows that the outer part of the diagram in Figure 6.10(b) commutes. Hence, Z is a Λ-bisimulation.
Item 2: Let $Z \subseteq X \times Y$ be a precocongruence relation with pushout (P, p_l, p_r), and let (U, V) be Z-coherent. By Lemma 6.2.9, there is a map $g\colon pb(Q\pi_l, Q\pi_r) \to QP$ such that $Qp_l \circ g = \overline{\pi}_l$ and $Qp_r \circ g = \overline{\pi}_r$. Then, by naturality of $[\![\heartsuit]\!]$ and the fact that p_l and p_r are T-coalgebra morphisms, it follows that the outer part of the diagram in Figure 6.10(b) commutes. Hence, Z is a Λ-bisimulation. $\square$

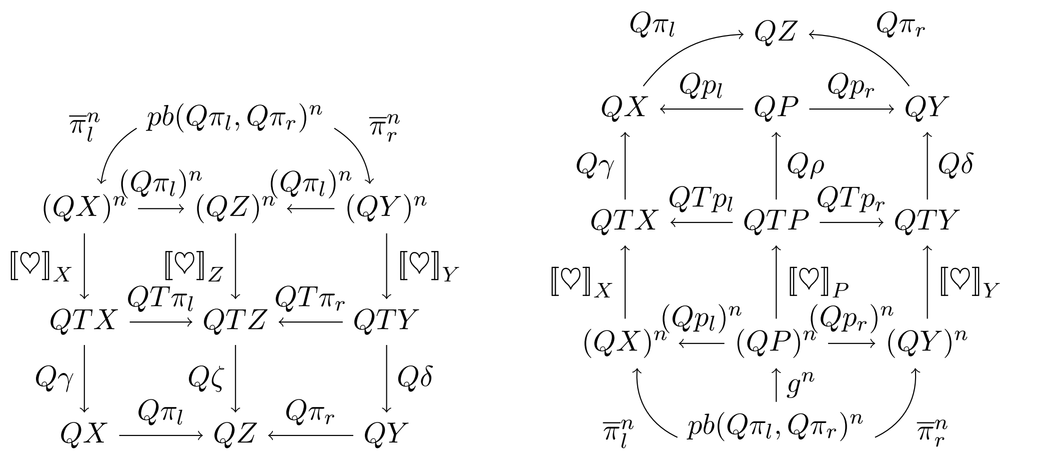

(a) T-bisimulations are Λ-bisimulations. (b) Precocongruences are Λ-bisimulations.

Figure 6.10. Precongruences and T-bisimulations are Λ-bisimulations.

The next proposition shows that, if Λ is separating, then we have the converse of Proposition 6.3.11(2).

Proposition 6.3.12 *If Λ is separating and $Z \subseteq X \times Y$ is a Λ-bisimulation between $\mathbb{X}$ and $\mathbb{Y}$, then Z is a precocongruence between $\mathbb{X}$ and $\mathbb{Y}$.* $\dashv$

Proof Let $Z \subseteq X \times Y$ be a Λ-bisimulation with projections $\pi_l : Z \to X$ and $\pi_r : Z \to Y$, and pushout (P, p_l, p_r) We need to define $\rho : P \to TP$ such that $\rho \circ p_l = Tp_l \circ \gamma$ and $\rho \circ p_r = Tp_r \circ \delta$. We obtain such a ρ from the universal property of the pushout, if we can show that for all $(x, y) \in Z$: $Tp_l(\gamma(x)) = Tp_r(\delta(y))$. To prove this, since Λ is separating, it suffices to show that for arbitrary $\heartsuit \in \Lambda$, n-ary, and $A_1, \ldots, A_n \subseteq P$, $Tp_l(\gamma(x)) \in [\![\heartsuit]\!]_P(A_1, \ldots, A_n)$ iff $Tp_r(\delta(y)) \in [\![\heartsuit]\!]_P(A_1, \ldots, A_n)$, which is equivalent to, $Q\pi_l \circ Q\gamma \circ QTp_l \circ [\![\heartsuit]\!]_P = Q\pi_r \circ Q\delta \circ QTp_r \circ [\![\heartsuit]\!]_P$. This holds because of the commutativity of the diagram in Figure 6.3.12, where the map h is obtained from Lemma 6.2.9. $\square$

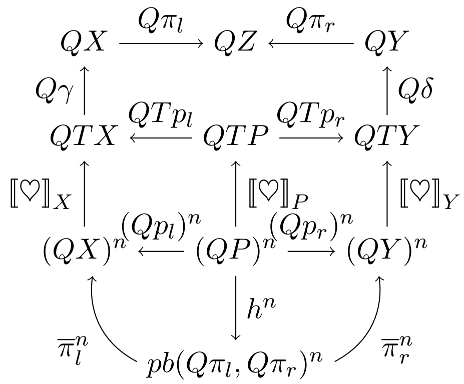

Figure 6.11. Proof of Proposition 6.3.12

It was shown in [89, Proposition 3.10] that, in general, T-bisimilarity implies precocongruence equivalence which in turn implies behavioural equivalence. This fact together with Proposition 6.3.12 tells us that Λ-bisimilarity implies behavioural equivalence, whenever Λ is separating. As mentioned in Subsection 6.2.4 if T preserves weak pulbacks, then T-bisimilarity coincides with behavioural equivalence. Hence in this case, by Proposition 6.3.12, it follows that Λ-bisimilarity coincides with T-bisimilarity and behavioural equivalence. The following proposition summarises our discussion so far.

Proposition 6.3.13 *Let Λ be a set of predicate liftings for T.*

1. $\mathbb{X}, x \sim_T \mathbb{Y}, y \quad \Longrightarrow \quad \mathbb{X}, x \sim_p \mathbb{Y}, y \quad \Longrightarrow \quad \mathbb{X}, x \sim_\Lambda \mathbb{Y}, y.$
2. *If Λ is separating, then*

$$\mathbb{X}, x \sim_p \mathbb{Y}, y \quad \Longleftrightarrow \quad \mathbb{X}, x \sim_\Lambda \mathbb{Y}, y \quad \Longrightarrow \quad \mathbb{X}, x \sim_{bh} \mathbb{Y}, y.$$

3. *If Λ is separating and T preserves weak pullbacks, then all four notions coincide:*

$$\mathbb{X}, x \sim_T \mathbb{Y}, y \quad \Longleftrightarrow \quad \mathbb{X}, x \sim_p \mathbb{Y}, y \quad \Longleftrightarrow \quad \mathbb{X}, x \sim_\Lambda \mathbb{Y}, y \quad \Longleftrightarrow \quad \mathbb{X}, x \sim_{bh} \mathbb{Y}, y.$$

⊣

The next lemma states that similar to the fact that T-coalgebra morphisms preserve and reflect behavioural equivalence, one can show that they preserve and reflect Λ-bisimilarity as well. We will use this fact to prove the Hennessy-Milner theorem in Section 6.4.

Proposition 6.3.14 *If $f : \mathbb{X} \to \mathbb{Y}$ is a T-coalgebra morphism, then for all $x, x' \in X$:*

$$\mathbb{X}, x \sim_\Lambda \mathbb{X}, x' \quad \textit{iff} \quad \mathbb{Y}, f(x) \sim_\Lambda \mathbb{Y}, f(x').$$

Proof For the direction from left to right, assume $\mathbb{X}, x \sim_\Lambda \mathbb{X}, x'$. Then, there exists a Λ-bisimulation Z on $\mathbb{X}$ such that $(x, x') \in Z$. We show that $(f \times f)(Z) = \{(f(x), f(x')) \in Y \times Y \mid (x, x') \in Z\}$ is a Λ-bisimulation. Let $(f(x), f(x')) \in (f \times f)(Z)$ and $\heartsuit \in \Lambda$. Note that if (U, V) is $(f \times f)(Z)$-coherent, then the pair $(f^{-1}[U], f^{-1}[V])$ is Z-coherent. By naturality and the fact that f is a coalgebra morphism, we have $\delta(f(x)) \in [\![\heartsuit]\!]_Y(U)$ iff $\gamma(x) \in [\![\heartsuit]\!]_X(f^{-1}[U])$, and $\delta(f(x')) \in [\![\heartsuit]\!]_Y(V)$ iff $\gamma(x') \in [\![\heartsuit]\!]_X(f^{-1}[V])$. Since Z is a Λ-bisimulation and $(f^{-1}[U], f^{-1}[V])$ is Z-coherent, we obtain $\delta(f(x)) \in [\![\heartsuit]\!]_Y(U)$ iff $\delta(f(x')) \in [\![\heartsuit]\!]_Y[V]$. A similar argument shows that if Z is a Λ-bisimulation on $\mathbb{Y}$ then $(f^{-1} \times f^{-1})(Z) = \{(x, x') \in X \times X \mid (f(x), f(x')) \in Z\}$ is a Λ-bisimulation on X. □

Λ-bisimulations: a different approach

Gorín and Schröder introduced in [79] a similar notion of Λ-bisimulation. To distinguish their notion from the one presented here, we refer to it as GS-Λ-bisimulation. One difference with our work is that Gorín and Schröder assume that Λ is a set of *monotone* predicate liftings. For convenience, we recall their definition. It can be stated without the assumption of monotonicity.

Definition 6.3.15 (GS-Λ-bisimulation) *A relation $Z \subseteq X \times Y$ is a GS-Λ-*bisimulation *if whenever $(x, y) \in Z$ then for all $\heartsuit \in \Lambda$ with $\heartsuit$ n-ary, and for all $A_1, \ldots, A_n \subseteq X$ and $B_1, \ldots, B_n \subseteq Y$*

(Zig-GS) $\gamma(x) \in [\![\heartsuit]\!]_X(A_1, \ldots, A_n) \Rightarrow \delta(y) \in [\![\heartsuit]\!]_Y(Z[A_1], \ldots, Z[A_n])$

(Zag-GS) $\delta(y) \in [\![\heartsuit]\!]_Y(B_1, \ldots, B_n) \Rightarrow \gamma(x) \in [\![\heartsuit]\!]_X(Z^{-1}[B_1], \ldots, Z^{-1}[B_n])$ ⊣

Under the assumption that all $\heartsuit \in \Lambda$ are monotone, it is straightforward to show that a GS-Λ-bisimulation is also a Λ-bisimulation. Example 6.3.5 demonstrates that there exists a choice of T and monotone Λ such that the two notions differ at the level of relations. Namely, the relation Z given there is a Λ-bisimulation, but not a GS-Λ-bisimulation. To see this, take $A = \{x_1, x_2\}$. We have that $\gamma(x) = \{x_1\} \subseteq A$, but $\delta(y) = \{y_1, y_2\} \not\subseteq Z[A] = \{y_1\}$. However, the next proposition shows that under the assumption that Λ is monotone, difunctional (also called zig-zag closed) Λ-bisimulations are GS-Λ-bisimulations. We first recall the relevant definition.

Definition 6.3.16 (Difunctional relation) *Let $R \subseteq X \times Y$ be a binary relation. We say R is* difunctional *if for all $x, u \in X$ and $y, w \in Y$, whenever $(x, y) \in R$, $(u, y) \in R$ and $(u, w) \in R$, then $(x, w) \in R$.* ⊣

Proposition 6.3.17 *Let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ be two T-coalgebras. If Λ is monotone, then difunctional Λ-bisimulation $Z \subseteq X \times Y$ is a GS-Λ-bisimulation.* ⊣

Proof Suppose that Z is a difunctional Λ-bisimulation, and let $A \subseteq X$ and $B \subseteq Y$. We will show that Z satisfies the (**Zig-**GS) and (**Zag-**GS) conditions. We only prove (**Zig-GS**), since (**Zag-GS**) can be proved in a similar manner. Let $\heartsuit \in \Lambda$, and let $(x, y) \in Z$ be such that $\gamma(x) \in [\![\heartsuit]\!]_X(A)$. Define $A' \subseteq X$ to be the set $A' = \{x \in X : Z[\{x\}] \subseteq Z[A]\}$. By monotonicity of Λ, it follows that $\gamma(x) \in [\![\heartsuit]\!]_X(A \cup A')$. As Z is difunctional, one can easily show that $(A \cup A', Z[A])$ is Z-coherent, and then since Z is a Λ-bisimulation, it follows that $\delta(y) \in [\![\heartsuit]\!]_X(Z[A])$. □

One can easily show that the relation $\sim_\Lambda$ between any two coalgebras is a difunctional. Hence, the two notions of GS-Λ-bisimilarity and Λ-bisimilarity coincide for monotone Λ. In [79, Theorem 26] it was shown that when Λ is separating and monotone, then GS-Λ-bisimilarity coincides with behavioural equivalence, and hence under these assumptions, Λ-bisimilarity coincides both with GS-Λ-bisimilarity and with behavioural equivalence.

Proposition 6.3.18 *If Λ is separating and monotone, then*

$$\mathbb{X}, x \sim_{GS\text{-}\Lambda} \mathbb{Y}, y \iff \mathbb{X}, x \sim_\Lambda \mathbb{Y}, y \iff \mathbb{X}, x \sim_{bh} \mathbb{Y}, y.$$ ⊣

We point out that our results on Λ-bisimulation do not require Λ to be monotone. Furthermore, our aims and results differ from those of [79] where the starting point was to investigate simulations between T-coalgebras. In this context, GS-Λ-bisimulations arose naturally as two-way simulations. The results in [79] focus on identifying conditions that ensure that GS-Λ-bisimilarity coincides with behavioural equivalence and/or T-bisimilarity. Our approach is to accept that the language is not expressive, and show that Λ-bisimilarity allows us to generalise several results that are known to hold for expressive languages.

Example 6.3.19 Consider INL from Example 6.2.19 (i.e. $T = \mathcal{P}\mathcal{P}$). Since $[\![\Box_n]\!]$ is monotone [19], it follows that $\Box_n$-bisimilarity coincides with GS-$\Box_n$-bisimilarity. The predicate lifting $[\![\Box_0]\!]_X(A) = \{N \in \mathcal{P}(\mathcal{P}(X)) \mid \exists U \in N : U \subseteq A\}$ is similar to the monotone neighbourhood modality (which is usually interpreted in $\mathcal{N}$-coalgebras). It is straightforward to prove that GS-$\Box_0$-bisimulations coincide with monotonic bisimulations (see e.g. [19]). ⊣

In the following proposition, we give a zig-zag characterisation of GS-$\Box_n$-bisimulation over $\mathcal{P}\mathcal{P}$-coalgebras. For all sets X and U, $U \subseteq_n X$ means that $U \subseteq X$ and $|U| = n$.

Proposition 6.3.20 *Let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ be two $\mathcal{P}\mathcal{P}$-coalgebras. For $n \geq 1$, $Z \subseteq X \times Y$ is a GS-$\Box_n$-bisimulation iff for all $(x, y) \in Z$:*

(Zig)$_n$ $\forall U \neq \emptyset \colon U \in \gamma(x) \Longrightarrow \forall U' \subseteq_n U\, \exists V \neq \emptyset \colon V \in \delta(y), V \subseteq Z[U]$ *and* $U' \subseteq Z^{-1}[V]$.

(Zag)$_n$ $\forall V \neq \emptyset \colon V \in \delta(y) \Longrightarrow \forall V' \subseteq_n V\, \exists U \neq \emptyset \colon U \in \gamma(x), U \subseteq Z^{-1}[V]$ *and* $V' \subseteq Z[U]$. ⊣

Proof We first prove the direction from left to right. Suppose that $n \geq 1$ and $Z \subseteq X \times Y$ is a GS-$\Box_n$-bisimulation. We will show that Z satisfies (**Zig**)$_n$ and (**Zag**)$_n$.
(**Zig**)$_n$: Let $(x, y) \in Z$, $U \neq \emptyset$, $U \in \gamma(x)$, and $U' \subseteq_n U$ with $U' = \{x_1, .., x_n\}$. By the definition of $[\![\Box_n]\!]$, $\gamma(x) \in [\![\Box_n]\!]_X(\{x_1\}, ..., \{x_n\}, U)$. Since Z is a GS-$\Box_n$-bisimulation, (**Zig**-GS) implies that $\delta(y) \in [\![\Box_n]\!]_Y(Z[\{x_1\}], ..., Z[\{x_n\}], Z[U])$. Again by the definition of $[\![\Box_n]\!]$, we obtain that there is a nonempty subset $V \subseteq Y$ with $V \in \delta(y)$ such that $V \subseteq Z[U]$ and $Z[\{x_i\}] \cap V \neq \emptyset$, for all $i \in \{1, ..., n\}$. This means that, for every $x_i \in U'$, there is $v_i \in V$ such that $x_i \in Z^{-1}[\{v_i\}]$ for $i = 1, \ldots, n$. Hence, $U' \subseteq Z^{-1}[V]$. (**Zag**)$_n$ can be proved in a similar way using the condition (**Zag**-GS).

For the other direction, suppose that Z satisfies the conditions (**Zig**)$_n$ and (**Zag**)$_n$. We show that Z satisfies (**Zig**-GS) and (**Zag**-GS).
For (**Zig**-GS) suppose $\heartsuit \in \Lambda$, $(x, y) \in Z$. Let the subsets $A_1, ..., A_n, A_{n+1} \subseteq X$ be such that $\gamma(x) \in [\![\Box_n]\!]_X(A_1, ..., A_{n+1})$. By the definition of $[\![\Box_n]\!]$, there is $U \subseteq A_{n+1}$ such that $U \cap A_i \neq \emptyset$, for all $i = 1, ..., n$, so $U \neq \emptyset$. For each $i \in \{1, ..., n\}$, choose $x_i \in U \cap A_i$, and let $U' = \{x_1, ..., x_n\}$. By (**Zig**)$_n$, there is a nonempty subset $V \in \delta(y)$ such that $V \subseteq Z[U]$ and $U' \subseteq Z^{-1}[V]$. This means that $x_i \in Z^{-1}[V]$, for all $i \in \{1, \ldots, n\}$, i.e., for every $x_i \in U'$, there is $y_i \in V$ such that $(x_i, y_i) \in Z$, for $i = 1, ..., n$, i.e., for all $i \in \{1, ..., n\} : Z[A_i] \cap V \neq \emptyset$. Since $U \subseteq A_{n+1}$, we get that $V \subseteq Z[A_{n+1}]$. Hence, $\delta(y) \in [\![\Box_n]\!]_X(Z[A_1], \ldots, Z[A_{n+1}])$. A similar argument works for (**Zag**-GS). □

6.3.3 Λ-morphisms

Given the fact that the graph of a T-coalgebra morphism is a T-bisimulation (cf. [141, Theorem 2.5.]), it is natural to define a Λ-*morphism from* $\mathbb{X}$ *to* $\mathbb{Y}$ to be a function $f \colon X \to Y$ such that the graph $Gr(f) = \{(x, f(x)) \mid x \in X\}$ is a Λ-bisimulation. It then follows from Proposition 6.3.11(1) that T-coalgebra morphisms are also Λ-morphisms.

Moreover, one can show that unlike Λ-bisimulations, Λ-homomorphisms are closed under composition. Therefore, T-coalgebras together with Λ-morphisms form a category.

In Enqvist [54], a weak notion of morphism for T-coalgebras was proposed which, like ours, is parametric in a set Λ of predicate liftings. To distinguish this notion from ours, we refer to it as E-Λ-morphism. We briefly recall the definition (which we state only for unary ♡, as is the case in [54]).

Definition 6.3.21 (*E*-Λ-morphism) *Given two T-coalgebras $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$, a function $f\colon X \to Y$ is an* E-Λ-morphism from $\mathbb{X}$ to $\mathbb{Y}$ *if for all $B \subseteq Y$, $x \in X$, and $\heartsuit \in \Lambda$:*

$$\delta(f(x)) \in [\![\heartsuit]\!]_Y(B) \text{ implies } \gamma(x) \in [\![\heartsuit]\!]_X(f^{-1}[B]). \tag{6.4}$$

Taking $Z = Gr(f)$, by Lemma 5.2.3(2) we know that a pair (U, V) is Z-coherent iff $U = f^{-1}[V] = Qf(V)$. It then follows that Λ-morphisms are E-Λ-morphisms. However, an E-Λ-morphism may not a Λ-morphism, as the following example shows.

Example 6.3.22 Consider the following $\mathcal{P}$-coalgebras $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$, where $X = \{x, x_1, x_2\}$, $\gamma(x) = \{x_1, x_2\}$, $\gamma(x_1) = \gamma(x_2) = \emptyset$, $Y = \{y, y_1, y_2\}$, $\delta(y) = \{y_1, y_2\}$, and $\delta(y_1) = \delta(y_2) = \emptyset$. It is straightforward to check that the $f : X \to Y$ with $f(x) = y$, and $f(x_1) = f(x_2) = y_1$ is an E-Λ-morphism. But it is not a Λ-morphism because the (**Coherence**) condition fails for the $Gr(f)$-coherent pair $(\{x_1, x_2\}, \{y_1\})$, that is $\gamma(x) \subseteq \{x_1, x_2\}$, whereas $\delta(y) \not\subseteq \{y_1\}$. ⊣

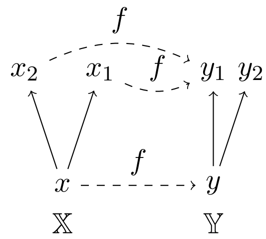

We do not investigate our notion of Λ-morphisms further in the present chapter. Several interesting questions could be asked, though. We discuss those in Section 6.5.

6.4 Hennessy-Milner theorem

This section is devoted proving a coalgebraic Hennessy-Milner theorem for Λ-bisimilarity. It is the main technical result of the chapter:

As we saw in Proposition 6.3.10, $\mathcal{L}_\Lambda$-formulas are invariant under Λ-bisimulation. Given that our modal language has only finite conjunctions, we will need to assume our coalgebra functor is finitary (see Def. 6.2.1). This is the analogue of restricting to image-finite Kripke frames (i.e., $\mathcal{P}_\omega$-coalgebras), as is done in the classic Hennessy-Milner theorem. As mentioned earlier, similar to the Hennessy-Milner for Δ-bisimilarity (Theorem 5.4.16), we define Hennessy-Milner classes of T-coalgebras with respect to $\sim_{\Lambda_+}$.

Definition 6.4.1 *A class* **C** *of T-coalgebras is a* Hennessy-Milner class, *if for every $\mathbb{X}$ and $\mathbb{Y}$ in* **C**, *we have $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$ iff $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{Y}, y$.* ⊣

As a first step towards our main result, we show that the class of finite T-coalgebras is a Hennessy-Milner class. We will use the following terminology. Given a T-coalgebra $\mathbb{X} = (X, \gamma)$, a subset $U \subseteq X$ is *modally coherent* if U is $\equiv_\Lambda$-closed. (Recall that $\equiv_\Lambda$ denotes the modal equivalence relation.) The next lemma provides us with a characterisation of modally coherent sets.

Lemma 6.4.2 *Let $\mathbb{X}$ be a finite T-coalgebra. For all $U \subseteq X$, U is modally coherent iff U is definable by a modal $\mathcal{L}_\Lambda$-formula.* ⊣

Proof It can be proved using the same line of argumentation as in the proof of Lemma 5.4.15(2). If $U = [\![\varphi]\!]_{\mathbb{X}}$ for some $\varphi \in \mathcal{L}_\Lambda$, then clearly U is modally coherent. For the converse implication, assume U is modally coherent, i.e., U is a union of modal equivalence classes: $U = \bigcup_{i \in I}[x_i]_{\equiv_\Lambda}$. Since X is finite, we may assume that I is finite. For $i, j \in I$ and $i \neq j$, there is a modal $\mathcal{L}_\Lambda$-formula $\delta_{i,j}$ such that $x_i \models \delta_{i,j}$ and $x_j \models \neg\delta_{i,j}$, so by taking $D_i = \{\delta_{i,j} \mid j \in I, i \neq j\}$, and using that I is finite, defining $\delta_i = \bigwedge D_i$, then we have that $[x_i] = [\![\delta_i]\!]_{\mathbb{X}}$ and $U = \bigcup_i [\![\delta_i]\!]_{\mathbb{X}}$. Therefore, U is definable by the formula $\delta = \bigvee_i \delta_i$. ☐

First, we state the finite version of Hennessy-Milner theorem for Λ-bisimulation.

Theorem 6.4.3 *Let $\mathbb{X} = (X, \gamma)$ and $\mathbb{Y} = (Y, \delta)$ be finite T-coalgebras, and let Λ be a set of predicate liftings for T.*

1. *For all states $x, x' \in X$: $\mathbb{X}, x \equiv_\Lambda \mathbb{X}, x'$ iff $\mathbb{X}, x \sim_\Lambda \mathbb{X}, x'$.*
2. *For all states $x \in X$ and $y \in Y$: $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$ iff $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{Y}, y$.* ⊣

Proof *Item 1*: The direction from right to left has been shown in Proposition 6.3.10. For the other direction, we show that $\equiv_\Lambda$ is a Λ-bisimulation. Let $x, x' \in X$ be such that $\mathbb{X}, x \equiv_\Lambda \mathbb{X}, x'$, and let $\heartsuit \in \Lambda$. For simplicity, we just give the argument for unary $[\![\heartsuit]\!]$. The n-ary generalisation is straightforward. Let $U \subseteq X$ be modally coherent. By Lemma 6.4.2, U is definable by a $\mathcal{L}_\Lambda$-formula ψ. We therefore have $x \in [\![\heartsuit\psi]\!]_{\mathbb{X}}$ iff $x' \in [\![\heartsuit\psi]\!]_{\mathbb{X}}$ because x and x' are modally equivalent. It follows that $\gamma(x) \in [\![\heartsuit]\!]_X(U)$ iff $\gamma(x') \in [\![\heartsuit]\!]_X(U)$. Hence, $\equiv_\Lambda$ is a Λ-bisimulation on $\mathbb{X}$.

Item 2: Follows from item 1 and the fact that the inclusion maps preserve truth of modal formulas: $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{Y}, y$ iff $\mathbb{X}+\mathbb{Y}, \mathsf{in}_l(x) \sim_\Lambda \mathbb{X}+\mathbb{Y}, \mathsf{in}_r(y)$ iff $\mathbb{X}+\mathbb{Y}, \mathsf{in}_l(x) \equiv_\Lambda$ iff $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$. ☐

We leverage the result for finite T-coalgebras to coalgebras for finitary functors.

Theorem 6.4.4 (Finitary Hennessy-Milner theorem) *Suppose T is a finitary functor, and $\mathbb{X} = (X, \gamma)$, $\mathbb{Y} = (Y, \delta)$ are T-coalgebras.*

1. *For all states $x, x' \in X$: $\mathbb{X}, x \equiv_\Lambda \mathbb{X}, x'$ iff $\mathbb{X}, x \sim_\Lambda \mathbb{X}, x.'$*
2. *For every $x \in X$ and $y \in Y$: $\mathbb{X}, x \equiv_\Lambda \mathbb{Y}, y$ iff $\mathbb{X}, x \sim_{\Lambda_+} \mathbb{Y}, y$.* ⊣

Proof *Item 1:* Let $x, x' \in X$ be such that $\mathbb{X}, x \equiv_\Lambda \mathbb{X}, x'$. By [3, Theorem 4.1] there exists a finite sub-coalgebra $\mathbb{X}_0 = (X_0, \gamma_0)$ of $\mathbb{X}$ with $x, x' \in X_0$. Since, the inclusion $\mathsf{in}_{X_0} : X_0 \to X$ is a T-coalgebra morphism and hence preserves truth of formulas, it follows that $\mathbb{X}_0, x \equiv_\Lambda \mathbb{X}_0, x'$. By Theorem 6.4.3(1) we obtain $\mathbb{X}_0, x \sim_\Lambda \mathbb{X}_0, x'$, and from Proposition 6.3.14, using again that in_{X_0} is a T-coalgebra morphism that $\mathbb{X}, x \sim_\Lambda \mathbb{X}, x'$.
Item 2: can be proved using item 1 in a similar way as item 2 of Theorem 6.4.3. ☐

6.5 Discussion and future work

We have defined a notion of Λ-bisimulation for weakly expressive coalgebraic modal logics, which is parametric in a collection of predicate liftings. The notions of rel-Δ-bisimulations and nbh-Δ-bisimulations of Chapter 5 are special cases of Λ-bismulations. We also have shown that our notion of Λ-bisimulation gives rise to a Hennessy-Milner theorem, which generalises the

Hennessy-Milner theorem for nbh-Δ-bisimilarity (Theorem 5.4.16). In other words, we have shown that Λ-bisimilarity fits exactly the expressiveness of the modal language.

As we discussed in Chapter 5, the coherence condition in the definition of Λ-bisimulation is, however, a non-local property as one would need to compute all coherent pairs over the state space in order to verify that two states are Λ-bisimilar. For concrete instances of Λ-bisimulations, it would be desirable to have a local back-and-forth style characterisation, similar to, e.g., the usual ones for Kripke frames, and the zig-zag conditions for Δ-bisimulations over Kripke frames in [62]. Such a local condition would obtain if Λ-bisimilarity could be charaterised in terms of relation liftings. In the case that Λ is separating, respectively monotone, Λ-bisimilarity coincides with precocongruences, respectively GS-Λ-bisimilarity, both of which have a relation lifting characterisation, cf. [89, 79]. We would like to investigate whether approaches such as those of [109, 117] can be used to obtain a relation lifting characterisation of Λ-bisimilarity under weaker conditions.

In Chapter 5, a Van Benthem characterisation theorem was proved for contingency logic over neighbourhood frames. That is, over neighbourhood frames, contingency logic is the fragment of first-order logic which is invariant under Λ-bisimilarity, where $\Lambda = \{\Delta\}$. We would like to generalise this result and show a coalgebraic version for Λ-bisimilarity, using as correspondence language *coalgebraic predicate logic (CPL)*, which was introduced in [111] as a first order correspondence language of coalgebraic modal logic.

We hardly explored the notion of Λ-morphisms in the present chapter. It would be interesting to know which constructions are possible in the category of T-coalgebras and Λ-morphisms. For example, in Chapter 5 we showed that for $T = \mathcal{N}$ and $\Lambda = \{\Delta\}$ one can construct Λ-quotients, i.e., quotients of T-coalgebras with respect to Λ-bisimilarity. We would like to know whether this is possible, in general. That would mean that we can minimise T-coalgebras with respect to Λ-bisimilarity. Finally, we would also like to know if a final object can be constructed from satisfied theories using techniques along the lines of [102, 125], and whether the Hennessy-Milner theorem for Λ-bisimilarity fits into the more abstract picture where a coalgebraic modal logic is obtained via a dual adjunctions, as in e.g. [99, 95].

References

[1] P. Aczel and N. Mendler. A final coalgebra theorem. In *Proceedings of Category Theory and Computer Science*, volume 389 of *Lecture Notes in Computer Science*, pages 357–365. Springer, 1989.

[2] J. Adámek, H. Herrlich, and G. E. Strecker. *Abstract and Concrete Categories - The Joy of Cats*. Dover Publications, 2009.

[3] J. Adámek and H. E. Porst. From varieties of algebras to covarieties of coalgebras. *Electronic Notes in Theoretical Computer Science*, 44(1):27–46, 2001.

[4] T. Ågotnes and H. van Ditmarsch. What will they say?—Public announcement games. *Synthese*, 179:57–85, 2011.

[5] O. Arieli and A. Avron. Reasoning with logical bilattices. *Journal of Logic, Language and Information*, 5(1):25–63, 1996.

[6] O. Arieli and A. Avron. The value of the four values. *Artificial Intelligence*, 102(1):97–141, 1998.

[7] R. J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4:1236–1239, 1976.

[8] S. Awodey. *Category Theory*. Oxford Logic Guides. Ebsco Publishing, 2006.

[9] A. Baltag. A coalgebraic semantics for epistemic programs. *Electronic Notes in Theoretical Computer Science*, 82(1):17–38, 2003.

[10] A. Baltag, B. Coecke, and M. Sadrzadeh. Algebra and sequent calculus for epistemic actions. *Electronic Notes in Theoretical Computer Science*, 126:27–52, 2005.

[11] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, pages 43–56. Morgan Kaufmann, 1998.

[12] S. Barbera, A. Bogomolnaia, and H. van der Stel. Strategy-proof probabilistic rules for expected utility maximizers. *Mathematical Social Sciences*, 35(2):89–103, 1998.

[13] J. J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[14] J. Barwise and J. Perry. *Situations and Attitudes*. Bradford books. MIT Press, 1985.

[15] N. D. Belnap. How a computer should think. In G. Ryle, editor, *Contemporary Aspects of Philosophy*. Oriel Press, 1977.

[16] N. D. Belnap. *A Useful Four-Valued Logic*, pages 5–37. Springer Netherlands, 1977.

[17] J. van Benthem. *Modal Correspondence Theory*. PhD thesis, Mathematisch Instituut and Instituut voor Grondslagenonderzoek, University of Amsterdam, 1977.

[18] J. van Benthem. Games in dynamic-epistemic logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.

[19] J. van Benthem, N. Bezhanishvili, S. Enqvist, and J. Yu. Instantial neighbourhood logic. *The Review of Symbolic Logic*, 10(1):116–144, 2017.

[20] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.

[21] D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.

[22] E. W. Beth. On padoa's method in the theory of definition. *Journal of Symbolic Logic*, 21(2):194–195, 1956.

[23] M. Bílková and M. Dostál. Expressivity of many-valued modal logics, coalgebraically. In *Proceedings of the 23rd International Workshop on the Logic, Language, Information, and Computation (WoLLIC 2016)*, volume 9803 of *Lecture Notes in Computer Science*, pages 109–124. Springer, 2016.

[24] M. Bílková, A. Palmigiano, and Y. Venema. Proof systems for the coalgebraic cover modality. In *Advances in Modal Logic 7*, pages 1–21. College Publications, 2008.

[25] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

[26] P. Blackburn, J. F. van Benthem, and F. Wolter. *Handbook of modal logic*, volume 3. Elsevier, 2006.

[27] J. C. de Borda. Mémoire sur les élections au scrutin. 1781.

[28] F. Bou, F. Esteva, L. Godo, and R. O. Rodríguez. On the minimum many-valued modal logic over a finite residuated lattice. *Journal of Logic and Computation*, 21(5):739–790, 2009.

[29] F. Bou and U. Rivieccio. Bilattices with implications. *Studia Logica*, 101(4):651–675, 2013.

[30] C. Boutilier and J. Rosenschein. Incomplete information and communication in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 10. Cambridge University Press, 2015.

[31] L. Bozzelli, H. van Ditmarsch, T. French, J. Hales, and S. Pinchinat. Refinement modal logic. *Information and Computation*, 239:303–339, 2014.

[32] F. Brandt, V. Conitzer, U. Endriss, A. D. Procaccia, and J. Lang. *Handbook of computational social choice*. Cambridge University Press, 2016.

[33] A. P. Brogan. Aristotle's logic of statements about contingency. *Mind*, 76(301):49–61, 1967.

[34] S. Burris and H.P. Sankappanavar. *A Course in Universal Algebra*. Graduate Texts in Mathematics. Springer New York, 2011.

[35] L. Cabrer, U. Rivieccio, and R. O. Rodríguez. Łukasiewicz public announcement logic. In *Proceedings of the 16th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU 2016)*, volume 611 of *Communications in Computer and Information Science*, pages 108–122. Springer, 2016.

[36] R. Carnap. Modalities and quantification. *The Journal of Symbolic Logic*, 11(02):33–64, 1946.

[37] B. F. Chellas. *Modal logic, an intoduction.* Cambridge University Press, 1980.

[38] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. volume 4362 of *Lecture Notes in Computer Science*, pages 51–69. Springer, 2007.

[39] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In *Proceedings of the 9th European Conference on Logics in Artificial Intelligence (JELIA 2004)*, volume 3229 of *Lecture Notes in Computer Science*, pages 18–30. Springer.

[40] C. Cirstea, A. Kurz, D. Pattinson, Schröder, and Y. Venema. Modal logics are coalgebraic. *The Computer Journal*, 54(1):31–41, 2008.

[41] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th Conference on Artificial Intelligence (AAAI 2011)*, volume 11, pages 638–643. AAAI Press, 2011.

[42] W. Conradie, S. Frittella, A. Palmigiano, and A. Tzimoulis. Probabilistic epistemic updates on algebras. In *Proceedings of the 5th International Workshop on Logic, Rationality, and Interaction (LORI 2015)*, volume 9394 of *Lecture Notes in Computer Science*, pages 64–76. Springer, 2015.

[43] W. Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *The Journal of Symbolic Logic*, 22(3):269–285, 1957.

[44] B. A. Davey and H. A. Priestley. *Lattices and Order.* Cambridge Univerity Press, 2002.

[45] H. van Ditmarsch and T. French. Simulation and information: Quantifying over epistemic events. In *Proceedings of the 1st International Workshop on Knowledge Representation for Agents and Multi-Agent Systems, (KRAMAS 2008)*, volume 5605 of *Lecture Notes in Computer Science*, pages 51–65. Springer, 2008.

[46] H. van Ditmarsch, T. French, and S. Pinchinat. Future event logic-axioms and complexity. *Advances in Modal Logic 8*, 8:77–99, 2010.

[47] H. van Ditmarsch, W. van der Hoek, and P. Iliev. Everything is Knowable – How to get to know whether a proposition is true. *Theoria*, 78(2):93–114, 2012.

[48] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Dynamic epistemic logic with assignment. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, pages 141–148. ACM, 2005.

[49] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic epistemic logic*, volume 337. Springer Science & Business Media, 2007.

[50] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Playing cards with Hintikka: An introduction to dynamic epistemic logic. *The Australasian Journal of Logic*, 3, 2005.

[51] D. Dubois. On ignorance and contradiction considered as truth-values. *Logic Journal of IGPL*, 16(2):195–216, 2008.

[52] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.

[53] E. Elkind, U. Grandi, F. Rossi, and A. Slinko. Gibbard-Satterthwaite games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 533–539. AAAI Press, 2015.

[54] S. Enqvist. Homomorphisms of coalgebras from predicate liftings. In *Proceedings of the 5th International Conference on Algebra and Coalgebra in Computer Science (CALCO 2013)*, volume 8089 of *Lecture Notes in Computer Science*, pages 126–140. Springer, 2013.

[55] D. W. Etherington. *Reasoning with Incomplete Information*. Morgan Kaufmann Publishers Inc., 1988.

[56] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *Journal of the ACM (JACM)*, 41(2):340–367, 1994.

[57] R. Fagin and J.Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39–76, 1988.

[58] R. Fagin, Y. Moses, J.Y. Halpern, and M.Y. Vardi. *Reasoning About Knowledge*. The MIT Press paperback series. MIT Press, 2003.

[59] J. Fan. *Logical studies for non-contingency operator*. PhD thesis, Peking University, 2015. (in Chinese).

[60] J. Fan. A note on non-contingency logic (manuscript), 2016.

[61] J. Fan and H. van Ditmarsch. Neighborhood contingency logic. In *Proceedings of the 6th Indian Conference on Logic and Its Applications (ICLA 2015)*, volume 8923 of *Lecture Notes in Computer Science*, pages 88–99. Springer, 2015.

[62] J. Fan, Y. Wang, and H. van Ditmarsch. Almost necessary. In *Advances in Modal Logic 10*, pages 178–196. College Publications, 2014.

[63] J. Fan, Y. Wang, and H. van Ditmarsch. Contingency and knowing whether. *The Review of Symbolic Logic*, 8(1):75–107, 2015.

[64] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.

[65] M. Fitting. Bilattices and the theory of truth. *Journal of Philosophical Logic*, 18(3):225–256, 1989.

[66] M. Fitting. Many-valued modal logics. *Fundamenta Informaticae*, 15(3-4):235–254, 1991.

[67] M. Fitting. Many-valued modal logics II. *Fundamenta Informaticae*, 17(1-2):55–73, 1992.

[68] M. Fitting. Bilattices are nice things. In *Proceedings of the PhiLog Conference on Self-Reference*. CSLI Publications, 2006.

[69] S. Frittella, G. Greco, A. Kurz, and A. Palmigiano. Multi-type display calculus for propositional dynamic logic. *Journal of Logic and Computation*, 26(6):2067–2104, 2016.

[70] S. Frittella, G. Greco, A. Kurz, A. Palmigiano, and V. Sikimić. Multi-type sequent calculi. In *Trends in Logic XIII: Gentzen's and Jaśkowski's Heritage 80 Years of Natural Deduction and Sequent Calculi*, pages 81–93. Lodź University Press, 2014.

[71] S. Frittella, G. Greco, A. Kurz, A. Palmigiano, and V. Sikimić. A proof-theoretic semantic analysis of dynamic epistemic logic. *Journal of Logic and Computation*, 2014.

[72] S. Frittella, G. Greco, A. Kurz, A. Palmigiano, and V. Sikimić. Multi-type display calculus for dynamic epistemic logic. *Journal of Logic and Computation*, 26(6):2017–2065, 2016.

[73] S. Frittella, G. Greco, A. Palmigiano, and F. Yang. A multi-type calculus for inquisitive logic. In *Proceedings of the 23rd International Workshop on the Logic, Language, Information, and Computation (WoLLIC 2016)*, volume 9803 of *Lecture Notes in Computer Science*, pages 215–233. Springer, 2016.

[74] D. M. Gabbay and L. M. *Interpolation and definability: modal and intuitionistic logics*, volume 1. Oxford University Press on Demand, 2005.

[75] J.D. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language, and Information*, 6:147–169, 1997.

[76] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.

[77] M. Ginsberg. Multivalued logics: A uniform approach to reasoning in artificial intelligence. *Computational intelligence*, 4(3):265–316, 1988.

[78] R. Goldblatt. Varieties of complex algebras. *Annals of Pure and Applied Logic*, 44(3):173–242, 1989.

[79] D. Gorín and L. Schröder. Simulations and bisimulations for coalgebraic modal logics. In *Proceedings of the 5th International Conference on Algebra and Coalgebra in Computer Science (CALCO 2013)*, volume 8089 of *Lecture Notes in Computer Science*, pages 253–266. Springer, 2013.

[80] G. Greco, A. Kurz, and A. Palmigiano. Dynamic epistemic logic displayed. In *Proceedings of the 4th International Workshop on Logic, Rationality, and Interaction (LORI 2013)*, volume 8196 of *Lecture Notes in Computer Science*, pages 135–148. Springer, 2013.

[81] J. Groenendijk and M. Stokhof. Questions. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 1055–1124. Elsevier, 1997.

[82] H. P. Gumm. State based systems are coalgebras. *Cubo–Matemática Educacional*, 5:239–262, 2003.

[83] J. Hales. Refinement quantifiers for logics of belief and knowledge. Honours Thesis, The University of Western Australia, 2011.

[84] J. Hales. Arbitrary action model logic and action model synthesis. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2013)*, pages 253–262. IEEE Computer Society, 2013.

[85] J. Hales. *Quantifying over epistemic updates.* PhD thesis, The University of Western Australia, 2016.

[86] J. Hales, T. French, and R. Davies. Refinement quantified logics of knowledge. *Electronic Notes in Theoretical Computer Science*, 278:85–98, 2011.

[87] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.

[88] J. Y. Halpern and L. D. Zuck. A little knowledge goes a long way: Knowledge-based derivations and correctness proofs for a family of protocols. *Journal of the ACM*, 39(3):449–478, 1992.

[89] H. H. Hansen, C. Kupke, and E. Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 5(2) (paper 2):1–38, 2009.

[90] J. C. Harsanyi. Games with incomplete information played by "bayesian" players, I-III: part i. The Basic Model&. *Management science*, 50(12-supplement):1804–1817, 2004.

[91] N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.

[92] A. Heifetz and P. Mongin. Probability logic for type spaces. *Games and Economic Behavior*, 35(1):31–53, 2001.

[93] M. Hennessy and R. Milner. Algebraic laws for non-determinism and concurrency. *Journal of the ACM*, 32:137–161, 1985.

[94] J. Hintikka. *Knowledge and Belief, An Introduction to the Logic of the Two Notions.* Cornell University Press Ithaca, 1962. Republished in 2005 by King's College, London.

[95] B. Jacobs and A. Sokolova. Exemplaric expressivity of modal logics. *Journal of logic and computation*, 20(5):1041–1068, 2010.

[96] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.

[97] R. Jansana and U. Rivieccio. Residuated bilattices. *Soft Computing*, 16(3):493–504, 2012.

[98] A. Jung and U. Rivieccio. Kripke semantics for modal bilattice logic. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2013)*, pages 438–447. IEEE Computer Society, 2013.

[99] B. Klin. Coalgebraic modal logic beyond sets. *Electronic Notes in Theoretical Computer Science*, 173:177–201, 2007.

[100] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of Multidisciplinary Workshop on Advances in Preference Handling (IJCAI 2005)*, volume 20, pages 124–129, 2005.

[101] S. A Kripke. Semantical analysis of modal logic in normal modal propositional calculi. *Mathematical Logic Quarterly*, 9(5-6):67–96, 1963.

[102] C. Kupke and R. A. Leal. Characterising behavioural equivalence: Three sides of one coin. In *Proceedings of the 3rd International Conference on Algebra and Coalgebra in Computer Science (CALCO 2009)*, volume 5728 of *Lecture Notes in Computer Science*, pages 97–112. Springer, 2009.

[103] C. Kupke and D. Pattinson. Coalgebraic semantics of modal logics: an overview. *Theoretical Computer Science*, 412(38):5070–5094, 2011.

[104] A. Kurz. *Logics for coalgebras and applications to computer science*. PhD thesis, Ludwig-Maximilians-Universität München, 2000.

[105] A. Kurz and A. Palmigiano. Epistemic updates on algebras. *Logical Methods in Computer Science*, 9(4:17):1–28, 2013.

[106] G. Lakemeyer. Tractable meta-reasoning in propositional logics of belief. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI 1987)*, volume 1, pages 402–408. Morgan Kaufmann, 1987.

[107] S. Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer, 1998.

[108] H. J. Levesque. A logic of implicit and explicit belief. In *Proceedings of the National Conference on Artificial Intelligence (AAAI 1984)*, pages 198–202. AAAI Press, 1984.

[109] P.B. Levy. Similarity quotients as final coalgebras. In *Proceeding of the 14th International Conference on Foundations of Software Science and Computational Structures (FoSSaCS 2011)*, volume 6604 of *Lecture Notes in Computer Science*, pages 27–41. Springer, 2011.

[110] K. Leyton-Brown and Y. Shoham. Essentials of game theory: A concise multidisciplinary introduction. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2(1):1–88, 2008.

[111] T. Litak, D. Pattinson, K. Sano, and L. Schröder. Coalgebraic predicate logic. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming (ICALP 2012)*, volume 7392 of *Lecture Notes in Computer Science*, pages 299–311. Springer, 2012.

[112] A. Lomuscio. *Knowledge sharing among ideal agents*. PhD thesis, University of Birmingham, 1999.

[113] J. Łukasiewicz. *Aristotle's syllogistic from the standpoint of modern formal logic*. Clarendon Press, 1957.

[114] M. Ma, A. Palmigiano, and M. Sadrzadeh. Algebraic semantics and model completeness for intuitionistic public announcement logic. *Annals of Pure and Applied Logic*, 165(4):963–995, 2014.

[115] L. Maksimova. The Beth properties, interpolation and amalgamability in varieties of modal algebras. In *Soviet Math. Dokl*, volume 44, pages 327–331, 1992.

[116] L. Maksimova. Definability and interpolation in classical modal logics. *Contemporary Mathematics*, 131:583–583, 1993.

[117] J. Marti and Y. Venema. Lax extensions of coalgebra functors and their logic. *Journal of Computer and System Sciences*, 81(5):880–900, 2015.

[118] M. J. Marx. *Algebraic relativization and arrow logic.* PhD thesis, University of Amsterdam, 1995.

[119] R. Meir. Plurality voting under uncertainty. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (IJCAI 2015)*, pages 2103–2109. AAAI Press, 2015.

[120] R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *ACM Conference on Economics and Computation, (EC 2014)*, pages 313–330, 2014.

[121] R. Montague. Universal grammar. *Theoria*, 36:373–398, 1970.

[122] H. Montgomery and R. Routley. Contingency and non-contingency bases for normal modal logics. *Logique et analyse*, 9(35/36):318–328, 1966.

[123] G. E. Moore. A reply to my critics. In P. A. Schilpp, editor, *The Philosophy of G. E. Moore*, pages 535–677. Northwestern University, Evanston IL, 1942.

[124] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96(1-3):277–317, 1999.

[125] L. S. Moss. Harsanyi type spaces and final coalgebras constructed from satisfied theories. *Electronic Notes in Theoretical Computer Science*, 106:279–295, 2004.

[126] R. Nelken and N. Francez. Bilattices and the semantics of natural language questions. *Linguistics and Philosophy*, 25(1):37–64, 2002.

[127] S. Odintsov and H. Wansing. Modal logics with belnapian truth values. *Journal of Applied Non-Classical Logics*, 20(3):279–304, 2010.

[128] M. J. Osborne and A. Rubinstein. *A course in game theory.* MIT press, 1994.

[129] E. Pacuit. Some comments on history based structures. *Journal of Applied Logic*, 5(4):613–624, 2007.

[130] R. Parikh, C. Tasdemir, and A. Witzel. The power of knowledge in games. *IGTR*, 15(4):1340030, 2013.

[131] D. Pattinson. Coalgebraic modal logic: soundness, completeness, and decidability of local consequence. *Theoretical Computer Science*, 309(1-3):177–193, 2003.

[132] D. Pattinson. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame Journal of Formal Logic*, 45(1):19–33, 2004.

[133] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.

[134] C. Pizzi. Contingency logics and propositional quantification. *Manuscrito*, 22(2):283, 1999.

[135] J. A. Plaza. Logics of public communications. In *ISMIS 1989*, pages 201–216. Oak Ridge National Laboratory, 1989.

[136] H. Reichgelt. Logics for reasoning about knowledge and belief. *The Knowledge Engineering Review*, 4(2):119–139, 1989.

[137] U. Rivieccio. *An algebraic study of bilattice-based logics.* PhD thesis, University of Barcelona, 2010.

[138] U. Rivieccio. Algebraic semantics for bilattice public announcement logic. In *Trends in Logic XIII*, pages 199–215. Lodz University Press, 2014.

[139] U. Rivieccio. Bilattice public announcement logic. *Advances in Modal Logic*, 10:459–477, 2014.

[140] U. Rivieccio, A. Jung, and R. Jansana. Four-valued modal logic: Kripke semantics and duality. *Journal of Logic and Computation*, 27(1):155–199, 2017.

[141] J. J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249(1):3–80, 2000.

[142] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[143] L. Schröder. Expressivity of coalgebraic modal logic: The limits and beyond. *Theoretical Computer Science*, 390(2-3):230–247, 2008.

[144] L. Schröder and D. Pattinson. PSPACE bounds for rank-1 modal logics. *ACM Trans. Comput. Log.*, 10(2):13:1–13:33, 2009.

[145] Lutz Schröder and Dirk Pattinson. Strong completeness of coalgebraic modal logics. In *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, (STACS 2009)*, volume 3 of *LIPIcs*, pages 673–684, 2009.

[146] D. Scott. Advice on modal logic. In *Philosophical problems in logic*, pages 143–173. Springer, 1970.

[147] F. Seifan, L. Schröder, and D. Pattinson. Uniform interpolation for coalgebraic modal logic. In *the 7th Conference on Algebra and Coalgebra in Computer Science (CALCO 2017)*. Springer, 2017. to appear.

[148] K. Sim. Epistemic logic and logical omniscience: A survey. *International Journal of Intelligent Systems*, 12(1):57–81, 1997.

[149] K. M. Sim. Beliefs and bilattices. In *ISMIS 1994*, pages 594–603. Springer, 1994.

[150] A. Slinko and Sh. White. Is it ever safe to vote strategically? *Social Choice and Welfare*, 43(2):403–427, 2014.

[151] C. Steinsvold. A note on logics of ignorance and borders. *Notre Dame J. Formal Logic*, 49(4):385–392, 2008.

[152] S. K. Thomason. Categories of frames for modal logics. *The Journal of Symbolic Logic*, 40(3):439–442, 1975.

[153] W. van der Hoek and A. Lomuscio. A logic for ignorance. *Electronic Notes in Theoretical Computer Science*, 85(2)(2):117–133, 2004.

[154] M. Y. Vardi. On epistemic logic and logical omniscience. In *Proceedings of the 1st Conference on Theoretical Aspects of Reasoning about Knowledge (TARK 1986)*, pages 293–305. Morgan Kaufmann, 1986.

[155] G. H. von Wright. *An essay in modal logic*. Norh-Holland, 1951.

[156] Y. Wang. Beyond knowing that: A new generation of epistemic logics. In H. van Ditmarsch and G. Sandu, editors, *Jaakko Hintikka on knowledge and game theoretical semantics*, Outstanding contributions to logic. Springer, 2016.

[157] Y. Wang and Q. Cao. On axiomatizations of public announcement logic. *Synthese*, 190(1):103–134, 2013.

[158] L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.