
NORMATIVE CHANGE: AN AGM APPROACH

JULIANO S. A. MARANHÃO
University of São Paulo, Brazil
julianomaranhao@usp.br

GIOVANNI CASINI
CNR-ISTI, Italy, and University of Cape Town, South Africa
giovanni.casini@isti.cnr.it

LEENDERT VAN DER TORRE
University of Luxembourg
leon.vandertorre@uni.lu

GABRIELLA PIGOZZI
Université Paris-Dauphine, France
gabriella.pigozzi@dauphine.fr

Abstract

Studying normative change is of practical and theoretical interest. Changing legal rules pose interpretation problems in determining the content of legal rules. The question of interpretation is tightly linked to questions about determining the validity of rules and their ability to produce effects. Different formal models of normative change seem to be better suited to capturing these different dimensions: the dimension of validity appears to be better captured by the AGM approach, while syntactic methods are better suited to modelling how the effects of rules are blocked or enabled. Historically, the AGM approach to belief revision (on which we focus in this article) was the first formal model of normative change. We provide a survey of the AGM approach along with the main criticisms of it. We then turn to a formal analysis of normative change that combines AGM theory and input/output logic, thereby allowing a clear distinction between norms and obligations. Our approach addresses some of the difficulties of normative change, like combining constitutive and regulative rules (and the normative conflicts that may arise from such a combination), revision and contraction of normative systems, as well as contraction of normative systems that combine sets of constitutive and regulative rules. We end our paper by highlighting and discussing some challenges and open problems with the AGM approach regarding normative change.

1 Normative Change and Legal Reasoning

The study of normative change in identifying the law and understanding legal reasoning and legal interpretation is of practical and theoretical interest.

From a practical perspective, legal rules are the product of, or at least affected by, the continuous agency of authorities with the power to issue norms or make judicial decisions.¹ Such authoritative acts change the content of the normative order by including and excluding rules or by modifying their effects.

The problem lies in the fact that there are a variety of acts that perform such modifications in the lifetime of a normative system, which may have an effect on two dimensions:

- (i) **validity**: the pertinence of rules to a normative system that may be changed by acts of abrogation, explicit derogation or implicit derogation;
- (ii) **efficacy**: the capacity of rules to produce effects or apply in a certain time period, which may be changed by acts of annulment or invalidation, suspension, restriction, modulation etc.

Hence, there is a discrepancy between the period of the validity of a rule in a normative system (which also has its own time span of existence), and its period of efficacy, thus creating situations where a rule is invalid but applicable or where a rule is valid but inapplicable.

From a theoretical perspective, it is important to understand normative change in order to understand the status of entailed (derived) rules in a normative system and their relationship to explicitly promulgated rules. The debate about the status of entailed rules is connected to a central problem in the conception of modern law concerning the role of reason *versus* the role of authority in identifying the law [51]. The question is whether the ultimate basis for identifying the legal status of an action are considerations of moral correction or goodness, or determination by a social source, i.e. whether the legal status of an action is determined by the content of an authoritative act, which is objectively identifiable independently of moral or political arguments [62].

¹Even scholars like Dworkin [23] who refuse to reduce identifying the law to the content of authoritative social sources do acknowledge that those sources produce relevant legal material for legal interpretation, potentially affecting how the law is identified and causing modification to the law.

1.1 Normative Change and Legal Validity

The inclusion of a new rule in a normative system is performed by an act of promulgation (or enactment). This new rule may represent new content, changing the content of the normative system by making new obligations, permissions or prohibitions derivable. Or the new rule may be redundant, adding a new norm-formulation, new text, without actually introducing new content.

In turn, exclusion of a rule from the normative system or modification of its effects may be obtained by means of a variety of legislative or judicial acts. There are terminological variations and disputes concerning acts that either exclude content pertaining to normative systems or restrict its efficacy (applicability). There are also different practices depending on the jurisdiction, and particularly with respect to systems of common law vis-à-vis systems of statutory law. In order to avoid confusion, we shall use terms in accordance with their technical usage in legal practice, but will articulate their meanings where the terminology can be misleading. In general, we will use the terms *derogation* and *abrogation* to refer strictly to the *dimension of validity*, with the meaning that a statute is totally or partially excluded from (ceases to pertain to) the normative system. We prefer to restrict the term “annulment” to the dimension of efficacy, with the meaning that a rule or a set of rules has its effects cancelled (ceases to be applicable).

Derogation is a distinct normative act that excludes a rule or some rules from a set of valid rules. It may be explicit or implicit:

explicit derogation: a new rule that explicitly mentions the name of the rule or rules to be excluded.

implicit derogation: a new rule that adds normative content which is inconsistent with the content of previous rules in the normative system.

In the case of explicit derogation, the content of the new rule may consist of only excluding the named rule: for instance, “article 56 of Law 1234 is derogated”. In such a case, the derogation rule exhausts its effects by performing that very derogation [38].

Abrogation means excluding the totality of the rules of a statute. Usually, abrogation is due to an act of promulgating a new statute that substitutes the content of a previous statute on the same subject. The exclusion is explicit because the set of excluded rules is indicated by either naming the statute or indicating the subject-matter. Abrogation also introduces new content whose effects hold after the previous statute has been derogated.

Derogation and abrogation (as well as promulgation) are usually non-retroactive normative acts, producing their effects immediately after publication or at a certain time in the future indicated by the same act. In the legal jargon, their effects are *ex nunc*, i.e. “from now on”. That is, they are “established” by the legislative act.

We shall use the term “annulment”² to refer to acts that cancel the effects of a valid rule. If a rule is annulled, it becomes inapplicable, that is, one cannot derive obligations, permissions, powers or any legal consequences from it.

An annulment may be the consequence of a judicial declaration that a rule of the normative system is invalid, or it may be the product of legislative acts cancelling the effects of a rule. A judicial annulment recognises a “vice” or “defect” in the “pedigree” of the rule. Those “pedigree” defects are related to problems with the source of the rule, the *legitimate authority*, the procedure for creating the rule, or the incompatibility of the rule with the content of hierarchically superior rules. Depending on the gravity of the defect identified, the recognition may consider the rule to be invalid from the time of its promulgation (in the legal jargon, *ex tunc* effects) or from the moment the defect is declared (*ex nunc*).

To complicate matters, since the annulment may be a judicial act, the recognition of invalidity may be general, that is, applicable to all legal subjects, or it may have an effect on a particular legal relation or a particular individual. So there is a general dimension of effects, but there are also indirect effects where normative changes affect the legal positions of different individuals in different ways. The same also happens for derogation and abrogation, which cannot retroact, so that a derogated rule may still be applicable to facts that occurred before the derogation took place.

There are other ways to affect the efficacy of rules by authoritative acts. A statute or decree may suspend or restrict the applicability of a rule in a given period or to a given domain or context. For instance, the legal rules protecting moral rights for authors became inapplicable to software by the force of a new law (art. 2 §1 of the Brazilian Law 9609/1998 on Software Copyright). Or a rule may suspend the applicability of some rental of real estate or labour laws during a global pandemic.

Clearly, the temporal aspect is crucial to analysing normative change, and this temporal factor has two dimensions: the time span of the rule’s validity, that is,

²The term “revocation” is sometimes used in parallel with annulment and pertains to the dimension of validity. Revocation refers to the act of cancelling a previous declaration, contract or legislative act, but the term “annulment” is also used to refer to such a cancellation with the intent of producing legal effects, particularly when such a cancellation is performed by a different person or institution (e.g. a judicial court) to the one that issued the act (e.g. the parliament or the contracting parties). Annulment and invalidation may also refer to cancelling the *effects* or applicability of a particular act, and are therefore situated in the dimension of efficacy of rules.

the period of time in which the rule pertains to the normative system; and the time span of its applicability, that is, the period of time where the obligations/permissions derived by the rule are applicable.

These dynamics of normative change, which are performed by a variety of legal acts with different effects, bring a series of difficulties for determining the content and the effects of a normative system at a particular moment in time. Indeed, a promulgation and a derogation may involve choices between alternative and incompatible descriptions of the resulting normative system.

The practical import of the study of normative change is not only a matter of finding suitable formal and computable representations of an uncontroversial and standard practice. It is also relevant for clarifying that very practice by describing the impact of acts of promulgation and revocation on the content of a normative system, and especially how they affect the normative consequences or entailed rules of that system. We highlight three problems.

The first problem concerns the *network effects of normative change*, that is, the effects of a derogation or a promulgation on networks of regulative and constitutive rules [71]. Acts of promulgation or derogation may not only add or exclude *regulative rules*, which are authoritative rules demanding, prohibiting or permitting an action or the omission of an action. They may also add or exclude *constitutive rules*, whose role is to a) define under which factual conditions a certain object or action “counts as” an instance of a legal concept such as property right, or b) ascribe meaning to legal concepts via definitions (e.g. people under 18 years old count as minors).

Hence, stipulating a new definition or changing the definition of a legal concept may affect how the content of different regulatory rules are determined. In turn, the exclusion or addition of new rules that are related to a legal concept may affect the practical implications, and therefore the very understanding, of that very concept [68]. Such an effect is neither immediately nor completely acknowledged by lawgivers, and leads to subsequent modifications and adaptations.

For instance, the legal definition of “software” as “literary work”³ makes rules protecting the “expression” of a literary work applicable to the source code of software: the copyright owner may copy, share, or distribute the software, create “derivative work” etc. The equiparation also enhances new legal consequences by analogy, such as the additional copyright protection of the original “structure” of a code, considering that the “composition” of different non-original literary works are also protected. Thus, the addition of new rules or protections for “literary work” may also “expand” the protection of software. However, some undesirable legal consequences of that equiparation—for instance, the ascription of “moral rights” related

³Agreement on Trade-Related Aspects of Intellectual Property Rights (Trips Treaty, 1994)

to software, such as the right to regret and withdraw the work from distribution—have been derogated in several jurisdictions.⁴ Such derogations in turn affect the understanding of the very concept of copyright—originally conceived as intrinsically bound to the author’s personality—by linking the original notion of copyright to a network of personality rights. Thus, the ascription of new objects to a legal concept by definitional rules and the introduction or derogation of regulatory rules interferes with, and demands “reconfigurations” of, the links in the network of legal definitions and normative consequences.

The second problem concerns the *undecidability of implicit derogations*, which is a consequence of the potential conflict between different “collision criteria” in the law. New obligations, prohibitions, permissions or definitions added via lawgiving acts may create conflicts with the content of the previous version of the normative system. Such conflicts are solved by an *implicit derogation* operated by so-called *collision criteria*, which are legal principles of interpretation enunciating preference relations for solving conflicts between rules. There are three collision criteria:

lex superior: a *hierarchical* criterion according to which rules enacted by a source of a higher hierarchical degree prevail over rules from lower degree sources.

lex posterior: a *temporal* criterion according to which more recent rules take precedence over older ones.

lex specialis: a criterion of *specialisation* according to which a rule applicable to a specific circumstance or condition prevails over another rule applicable in a more general context.

Although it is clear that the hierarchical criterion prevails over the temporal and speciality criteria, the two last criteria may collide.

Example 1. *Suppose that a new statute on public concessions is promulgated stating:*

1. *A private company operating a public concession of a federal road may explore its margins for commercial purposes.*

This rule might conflict with a previous existing rule specific to electricity distribution companies stating:

⁴For instance, article 2º, §1, of the Brazilian Copyright Law considers all provisions of the law warranting moral rights to be inapplicable to software, except for the right to have authorship acknowledged and the right to oppose unauthorised modifications that may affect the reputation of the author.

2. *Public energy distribution companies have the right to use road margins to the extent that such use is necessary to install its energy transmission network.*

These rules conflict if one interprets the right to use, in which energy companies are invested, as the right to use free of charge, and if the right to explore the margins “for commercial purposes” is considered to include a right to charge a fee for the public energy distribution system. But the conflict cannot be solved by the existing collision criteria because there is a conflict in this case between *lex posterior*, which makes rule (1) prevail over rule (2), and *lex specialis*, which makes rule (2) prevail over rule (1). Actually, there is another possible source of dispute, which is the understanding of which rule is the more specific rule. One could argue that rule (2) is more specific because it relates to a public energy distribution company, while rule (1) relates to all kinds of potential users. However, one could also argue that rule (1) is more specific because it relates to roads, the object of public concessions to private companies, while rule (2) has a wider scope on this aspect.

Hence, given a conflict of rules created by a promulgation, there may be no fixed criteria for deciding which one should prevail.

The third problem concerns the *indeterminacy of implicit derogations*, that is, that the promulgation of a new rule may conflict with a rule derived from the combination of different explicit rules in the normative system.

Example 2. *Suppose that a regulation contains the following rules:*

3. *Brasilia is the capital city of the Brazilian Federation.*
4. *The Brazilian Federal Administration must be located in the capital city of the Brazilian Federation.*

Now suppose that the following rule is promulgated:

5. *The Brazilian Federal Administration must be located in Rio de Janeiro.*

Rule 3 does not conflict with either rule 1 or 2, but it does conflict with the entailed rule:

- 5'. *The Brazilian Federal Administration must be located in Brasilia.*

This would be a case of *implicit derogation* of an entailed rule resolved by the temporal criteria of collision. However, the entailed rule can only be suppressed if at least one of explicit rules (3) or (4) are derogated. Hence, the content of the

normative system after the promulgation of (5) is undetermined, with three possible candidates for the outcome of this derogation: $S_1 = \{3, 5\}$, $S_2 = \{4, 5\}$ and $S_3 = \{5\}$.

From a domain-specific consideration, S_1 is plausible although it may have perplexing consequences (for instance, if there is a rule assigning a budget to the Brazilian capital that includes expenses for relocating and maintaining the offices of the Federal Administration). System S_2 would not properly imply that:

3'. Rio de Janeiro is the capital city of the Brazilian Federation.

But promulgating a norm specifying a city other than Rio de Janeiro as the capital city of Brazil would again lead to inconsistency.

Finally, system S_3 would leave the capital city of Brazil undefined, which could create uncertainty in the application of other rules employing that concept.

A similar problem of indeterminacy would appear when a rule entailed from a new and hierarchical superior rule is promulgated.

Example 3. *Suppose that a normative system contains the following rule:*

6. *All industries are free economic activities except for the public services listed below: (...)*

Suppose that the aviation industry is not listed in rule (6), implying that aviation is a free economic activity, and suppose also that there is a federal statute (the Aviation Code) stating the following:

7. *Aviation companies must be controlled by national investors.*

Now consider that a constitutional rule is enacted imposing the following:

8. *There ought to be no discrimination between the national and foreign capital of companies dedicated to any free economic activity.*

Considering that control by national investors counts as “discrimination” between foreign and national investors, rule (8) conflicts with rules (6) and (7), although originally the last two rules seemed to have no relevant connection to each other. The inconsistency is solved if either of these last two rules is derogated. The first option is to delete constitutive rule (6), which classifies the aviation industry as a free economic activity. The second option is to delete rule (7), thereby weakly permitting, that is not prohibiting, the control of aviation companies by foreign investors.

Hence, the interaction between constitutive and regulative rules, the problem of implicit derogation and the derogation of entailed rules all open up different possibilities for identifying the normative system resulting from normative revisions. Logical analysis of normative change should be faithful to such an indeterminacy, making the different possibilities for the resulting normative system transparent. Legal interpretation and argumentation may provide further constraints in order to select which, among all the possible candidates, would be the preferred outcome of a derogation, which may be domain-specific, or may have its rationality represented in formal models of normative change.

1.2 Normative Change and Legal Interpretation

Legal reasoning can be conceptually structured as three main tasks, as suggested by Wroblewski [77, 78]:

- (i) **validity**: identifying the valid legal rules that are generally applicable to the subject-matter;
- (ii) **interpretation**: determining the content of the rules identified as valid;
- (iii) **application**: instantiating the content of the valid rules applied to concrete or hypothetical cases (this last task includes identifying the relevant facts of the case, identifying how they qualify according to the applicable rules, and determining the legal consequences based on those rules).

At first glance, normative change should only be concerned with questions of validity, since the dynamics of promulgation and derogation determines the timeframe for the applicability of rules in normative systems. However, the three problems highlighted above show an intrinsic connection between normative change and legal interpretation, given that one of the main triggers of normative dynamics is the need to handle inconsistencies between the *content* of different rules in the normative system.

The problem of *network effects* is connected to determining the content of regulative rules with conceptual definitions. The *undecidability* problem is also about choosing between rules with conflicting content. The *indeterminacy* problem of implicit derogation concerns a conflict between the content of the promulgated rule and the content entailed by the normative system.

Given that the core task of legal interpretation is to determine the content of legal rules, it is necessary to first identify inconsistencies between rules, and therefore to check whether an implicit derogation has undermined the validity of a rule. Hence,

questions of validity and interpretation are not serial but circular. The object of interpretation is the content of valid rules, but interpretation is also necessary to the inquiry about validity. The same applies to interpretation and application. Since the conditions for applying the rule may not be isomorphic to the factors or circumstances of the case at hand [60, p. 77 ff.], the rule must be adapted to become “operational”. Further qualifications to the facts must be introduced via definitions that match the factual properties of the case with the concepts employed in the rule in order to make them isomorphic [1]. Hence, although it is the content of the rule that is subsequently instantiated, that instantiation induces modifications to the content of the rule to be applied [66, p. 36 ff.].

Hence, interpretation is pervasive in legal reasoning, performing an important role from identifying the authoritative sources to determining the legal effects on a concrete or hypothetical case.

Broadly understood, legal interpretation encompasses both *linguistic* and *constructive* interpretation. Linguistic interpretation consists in identifying the semantic/pragmatic content that is conveyed by an authoritative legal text.⁵ In turn, constructive interpretation, or “legal construction” [72], consists in determining the legal effect of that linguistic content, which means constructing the content of an “operational rule”.

Some conceive of linguistic interpretation as an inquiry into the linguistic facts of a language community [11, 72, 54], while others include an evaluative component in every linguistic inquiry [26, 23], and therefore consider the whole process of interpretation as constructing rules in the light of the purpose of legal practice. But even those who question the distinction accept that there would be a pre-interpretive stage where some preliminary meaning ascription takes place.

The linguistic interpretation or pre-interpretive stage may provide unsatisfactory solutions for a particular case. The linguistic meaning of the rule may not indicate a normative solution to a particular constellation of relevant facts [4], leaving a so-called “gap” in the normative system that must be fulfilled. The linguistic inquiry may also provide conflicting commands deriving from the same rule or from different rules, in which case the contradiction must be corrected. It may provide an array of alternative meanings (ambiguity), from which only one must be chosen, or may provide an imprecise meaning (vagueness), demanding further definitions to determine whether the case at hand fits the conditions for applying the rule. Finally, the rule’s command as determined by the linguistic inquiry may violate the rule’s underlying justification (the values promoted by the rule), which may necessitate

⁵Legal theorists disagree about what is the object of legal interpretation. While some contend that the object of interpretation is to formulate norms from authoritative sources [64], others, like Dworkin [23] would also include the whole argumentative social practice of law [22].

the introduction of exceptions or the specification of new conditions for applying the rule so that its content aligns with its purpose.

These further processes of

- filling gaps by adding new content,
- eliminating ambiguities by choosing between different content,
- eliminating vagueness by adding definitions to make the rule precise,
- resolving inconsistencies between rules by excluding content, and
- resolving deviances to the rule's command with respect to its underlying justification by modifying its conditions of application,

all clearly involve changes not only to the rule to be applied but also to the very normative system. The process of constructing an operational rule to be applied presupposes that the interpreted rule coheres with the normative system, and therefore that what is instantiated is actually a reconstructed fragment of a normative order containing a set of rules that are relevant to defining the deontic status (obligatory, forbidden, permitted) of the action at stake [4]. This reconstruction may be performed by a judge to solve a concrete case (judicial interpretation), or in legal doctrine when indicating solutions to hypothetical legal cases (doctrinal interpretation).

Note that in practice it is difficult to discriminate between these two different dimensions of legal interpretation—linguistic and constructive—considering that the very ascription of meaning to legal texts is constrained by a presumption of the lawgiver's rationality or "unity of will" [14], which requires that a text must be given a meaning that avoids inconsistencies or misalignments with the rule's purpose, and preferably avoids gaps and imprecision. Hence, construction may take place even when the identification of the meaning of a rule is uncontroversial.

For instance, consider the regulation on abortion in the Brazilian Criminal Code.

9. Causing an abortion; Punishment: imprisonment from 1 to 3 years.

10. Abortion performed by a physician is not punishable: (i) if there is no other way to save the pregnant woman's life; (ii) the pregnant woman has consented to the abortion and the pregnancy is the result of sexual abuse.

A criminal lawyer would say that it is settled from the text above that it is forbidden to abort if the pregnant woman's life is not endangered and no sexual abuse took place. Some would even say that this conclusion is immediate and does not

require interpretation. However, first of all, the interpretation of clauses (i) and (ii) as disjunctive and not conjunctive involves some evaluative considerations favouring women's freedom. Secondly, the plain language meaning actually reveals inconsistency between rules (9) and (10). Rule (10) is read as an exception, but this means that some interpretation cannons operate in order to first assume that inconsistent rules should be applicable to different hypothetical conditions, then to derogate (9) by specificity, and finally to reintroduce the prohibition of causing an abortion in scenarios that have not been exempted (*exceptio firmat regulam in casibus non exceptis*). The "operational rules" reconstructed from the original linguistic meaning are thus:

- 9*. Abortion is forbidden if not performed by a physician or if there are other ways to save the pregnant woman's life and the pregnancy is the result of sexual abuse or the pregnant woman has not consented to the abortion.
- 10*. Abortion is permitted if performed by a physician and there is no other way to save the pregnant woman's life or if the pregnancy is the result of sexual abuse and the pregnant woman has consented to the abortion.

The fact that what is assumed to be the "plain language meaning" of a norm already involves its construction leads some to consider the object of legal interpretation to be the legal community's set of settled instantiations of the valid rules [54] rather than the ordinary meaning of legal texts. In this conception, legal interpretation would then be the process of construction from that restricted basis of settled law, in order to develop solutions for unclear cases with gaps, imprecision and/or conflicts, etc.

Legal construction allows flexibility in the law so that it can adapt to new circumstances and social demands while reinforcing the authority of the normative order. It can achieve this by keeping track of the original rules (taking as a starting point the legal text, the clear and settled instantiations, or the legal history) and making them align with community values. Assessment of this interpretative practice from the perspective of normative change reveals different strategies used in legal doctrine, or by the courts, to manipulate the legal material in the sources in order to justify choosing a particular legal solution. Particularly interesting is their stipulation of definitions affecting relevant concepts of the rule.

Consider, for instance, the controversy in many jurisdictions concerning police access to the content of mobile phones in search & seizure orders.

In 2014, a decision by the Brazilian Superior Court of Justice (STJ: HC 51.531-RO) held that a WhatsApp conversation on a mobile phone collected in a search

procedure is analogous to ongoing correspondence and should count as “written communication”. Therefore, an *order to intercept* was mandatory to access its content, otherwise the access would have violated freedom of communication. However, in a decision reached in 2016 (STJ: HC 75.800-PR), the same court affirmed that a message exchange on a mobile phone is just stored data and therefore a property item which, according to the statutes, may be accessed in a search & seizure procedure.

The German Constitutional Court (BVerfGE, 115,166, *Kommunikationsverbindungsdaten*) also concluded that access to data stored on a mobile phone collected during an investigation does not violate rules regarding search & seizure. Such data would be analogous to information in a physical document since both involve possession and the data or information could have been destroyed by the searched individual. Therefore, accessing the history of calls does not affect freedom of communication, and does not have a greater impact on informational autonomy or property rights deserving special protection.

Example 4. *Consider a normative system with the following regulative rules:*

11. *Police officers have the power to access any property item if and only if authorised by a judicial search & seizure order.*
12. *Police officers have the power to intercept written or oral communication if and only if authorised by a judicial interception order.*

The following conceptual rules are key to determining whether stored text messages may be accessed in a search & seizure order:

13. *A message exchange stored on a mobile phone counts as ongoing communication;*
14. *A message exchange stored on a mobile phone counts as stored data;*
15. *Stored data counts as a property item.*

Suppose that officers only hold a search & seizure order. Then there is an inconsistency between conceptual rule (13), on the one hand, and conceptual rules (14) and (15) on the other. The difficulty lies in the fact that the linguistic meaning of a message exchange supports its qualification as both communication and stored data. The link between stored data and property pertains to the legal language and derives from valid legal rules. The German court has just excluded rule (13), thus avoiding that the search procedure should become unconstitutional by affecting freedom of communication. One of the Brazilian courts chose to delete rule (14).

But those qualifications (data as property, stored messages as data, message exchanges as communication) are also relevant to the application of other rules. Another solution to keep rule (11) compatible with the constitutional value of freedom of communication, and with a lower impact on the network of conceptual and regulative rules, would be to refine rule (11) as follows:

- 11*. Police officers have the power to access any property item, except for the digital content of mobile phones, if and only if authorised by a judicial search & seizure order.

Indeed, this was the solution adopted by the U.S. Supreme Court in a similar case involving search powers in an arrest (*Riley v. California*, 2014).

Hence, legal construction involves manipulating conceptual definitions not only by legal doctrine, but also regulative rules. This possibility does not offend the authority of the rules provided that, first, conceptual definitions may also be stipulated by valid legal rules, and secondly, that valid regulative rules may be derogated or refined by introducing exceptions, in the name of consistency with constitutional values, as explicit and higher order rules [8].

But it is clear that legal construction and legal interpretation in general have both a conservative and a creative component [22]. On the one hand, construction must be faithful to the settled normative order. On the other hand, it must enhance new solutions by clarifying the content of that order. In other words, choices and changes to the content of the legal order are going to take place, but only to the extent that is minimal and necessary to clarify its content.

It is also characteristic of such constructions that their conclusion is presented as entailing a coherent interpretation of the normative system. Opposing conclusions in apparently similar cases are shown to align with the balance of the relevant values pursued by the normative system. Alignment is attained by using an array of different techniques in constructive interpretation: discarding possible conceptual qualifications e.g. excluding the rule that stored messages count as communication, introducing exceptions to rules e.g. excluding mobile phones from the general search powers of officials, and introducing or excluding values from consideration.

It is clear from this discussion and examples that legal construction as a fundamental dimension of legal interpretation consists in making changes to the content of the normative system, and that these changes are driven by both a demand for coherence and by a demand for conservatism or “minimal change” to the legal order. These drivers show how logics of theory change are suitable for modelling legal construction.

To conclude this practical perspective, we observe that the relationship between interpretation and normative change is twofold. On the one hand, legal interpreta-

tion is a precondition to the dynamics of normative systems, as the identification of inconsistencies between the content of rules depends on it. On the other hand, the very activity of legal interpretation may be seen as dynamics of change affecting constitutive and regulatory rules.

1.3 Normative Change and Implied Rules

From a theoretical perspective, normative change is an important factor in understanding the status of implied (derived) rules in a normative system and its relation to explicitly promulgated rules. The debate about the status of entailed rules is connected to a central problem in the conception of modern law concerning the role of reason *versus* the role of authority in identifying the law. The question is whether the ground for identifying the legal status of an action consists in reasoning about its correction or goodness or whether this status is determined by the will of an authority with respect to individual or collective behaviour or its outcome.

If one conceives that the binding force of the content of explicit rules is the outcome of the authority's will manifested in the norm-giving act, the question arises whether or to what extent obligations, prohibitions or permissions deductively derived from those original rules, albeit not explicitly endorsed by the authority, are also binding or should also be considered to be part of the normative system.

This problem may be explored from the perspective of normative dynamics. Instead of a synchronic epistemology considering the identification of a rule as a matter of examining the foundational or coherentist connection of its content to the content of the other rules of the system [10], one may adopt a diachronic perspective of examining the vulnerability of the rule's content to changes in the normative system. If derived rules have the same "ontological status" as explicit rules, then, on the one hand, the promulgation (addition) of derived rules would be redundant and, on the other hand, their derogation would immediately mean a change in the normative system.

For instance, the Brazilian Criminal Code forbade sexual abuse with the following set of explicit rules:

16. It is forbidden to practice sexual intercourse without consent.
17. Sexual intercourse with a person under 14 years old shall be considered to be without consent.

Should we consider the derived rule (18) below a valid legal rule of the Brazilian criminal law system?

18. It is forbidden to practice sexual intercourse with a person under 14 years old.

A decade ago, a controversial decision by the Brazilian Supreme Court ruled that habeas corpus applied to an offender who maintained a sexual relationship with a 12 year old girl. The legal community has interpreted that ruling as *contra legem*, since it was widely assumed that the act violated the criminal code. It seems plain enough that although rule (18) was not explicitly promulgated, compliance with its content should be obligatory and any disregard would be a violation. And this follows from the fact that the content of (18) is deductively derived from rules (16) and (17).

Given that there is such a derived obligation, some would argue that rule (18) is also part of the normative system [4, 55]. Here, the binding force of the obligation is an outcome of reasoning (deduction), and if law is the system of binding rules, it should be part of the normative system as well.

Some, however, would accept the binding force of such derived rules, but would not acknowledge them as part of the normative system if their content is not explicitly willed [54]. Accepting them as part of the normative system, Marmor argues, would imply a (most probably) false assumption that the set of legal rules is coherent. Others, like Joseph Raz [63], would only accept them if such derivations were endorsed by the relevant authority (even though it is not quite clear what such endorsement means) as something distinct from explicitly willing its content but inferring such content from the explicit rules.

Curiously enough, that controversial decision by the Brazilian Supreme Court led to a legislative act (Law 12.015/2009) introducing rule (18) as an explicit rule of the Code. Did that law effectively change the Brazilian criminal law system? One could say that these are two different formulations of the Code representing the same criminal law system, provided that they contain the same set of derived obligations. If this is true, what led to the promulgation of the new legislative act?

One could say that it was fundamentally a political gesture with redundant or irrelevant legal consequences. Or one could say that the Supreme Court had actually changed the law, which was later modified by legislation again. But the interesting question is: if two different normative systems have identical normative consequences, is it the case that identical promulgations or derogations in each of these systems would lead to the same resulting normative system?

Example 5. *Consider normative system S1 with the following formulations:*

16. *It is forbidden to practice sexual intercourse without consent.*

19. *Sexual intercourse with a legally incompetent person shall be considered to be without consent.*

20. *A person becomes legally competent by reaching 14 years of age.*

Now consider normative system $S2$ containing rules (16), (19), (20) and, in addition, (18) as an explicit rule.

18. *It is forbidden to practice sexual intercourse with a person under 14 years old.*

Suppose now that the following rule is promulgated:

21. *A person becomes legally competent by reaching 16 years of age.*

Clearly, rule (18) is derived from $S1$. Hence, from the synchronic perspective, it is clear that $S1 = S2$, since the set of derived obligations is the same. But the effect of promulgating rule (21) in $S1$ is different from its promulgation in $S2$. In $S1$, promulgated rule (21) substitutes rule (20), and therefore the revised system ($S1^*$) derives the following:

(18*) It is forbidden to practice sexual intercourse with a person under 16 years old.

However, in system $S2$, rule (18) would still be derived. And while rules (20) and (21) conflict, this is not necessarily a conflict between explicit rule (18) and derived rule (18*). Therefore rule (18) could still be derivable. It would be a matter of legal interpretation to determine whether the new definition of legal competence would be applicable only to civil law, that is, the ability to perform valid civil and contractual acts, or whether it would also be applicable to criminal law, specifically, the ability to consent to sexual intercourse or to be liable to criminal responsibility.

Hence, from a synchronic perspective, i.e. considering the normative system at a particular moment in time, one may assume that two normative systems are the same if they derive the same set of obligations/permissions, even if they have different formulations. That is, from that perspective, the formulation of the base of explicit rules is irrelevant. However, from a diachronic perspective, that is, considering the normative system's change from one moment to a second moment where a new rule is promulgated or derogated, the formulation of the base of explicit rules becomes relevant, given that the revision of different sets of explicit rules with the same derived obligations/permissions may lead to different outcomes. Therefore, changes in the base of explicit rules may not result in changes in the set of obligations/permissions, but every change in the set of obligations/permissions means a change in the base.

This observation makes it clear that even if one assumes that the content of derived rules is as equally binding as the content of explicit rules, which would make these rules share the same “normative status”, it is not the case that they should share the same “pertinence status”. That is, the fact that a derived obligation is binding does not imply that it is a rule pertaining to the normative system.

1.4 Modelling Normative Change

The distinction between the dimension of the validity of a rule (the time span of the pertinence of a rule to the normative system) and the binding force or efficacy of derived obligations or permissions (the time span where obligations and permissions are applicable) is also relevant for defining an appropriate methodology for the study of normative change. The different methods may focus on one or another aspect of normative change, namely, changes to the content of norms that are part of the normative order, or changes with respect to the effectiveness of obligations over time.

Suppose that there is a normative system S_3 with the rule:

22. Abortion is forbidden.

Since this is an absolute prohibition, it applies to every possible circumstance. Therefore, the following prohibition is derived:

23. Abortion is forbidden if the pregnancy is the result of sexual abuse.

Suppose that a legislative or judicial authority wants to change rule (23) by permitting abortion in the case of sexual abuse (or a legal scholar argues that there is an “implicit exception” to the prohibition of abortion based on the constitutional value of a woman’s dignity). This normative change may be described in at least three different ways corresponding to three different methods proposed in the literature on artificial intelligence & law for modelling normative change.

The first methodology, devised by Governatori and Rotolo [30], may be called the *syntactic approach*. According to this approach, norm change is an operation performed on the rules contained in the code for determining whether a default rule is applicable or ceases to be applicable in defeasible deontic logic. So, the focus of the approach is not really the dimension of validity (the pertinence of the rule to the normative system) but the dimension of the efficacy (applicability) of derived obligations and permissions. They call “annulment” the operation where all the past and future effects of the rule are cancelled and “abrogation” the operation where

only the effects to the future are cancelled while past effects still hold. They use a temporal extension of defeasible logic to keep track of changes in the normative system and to deal with retroactivity (the possibility of changing the applicability of obligations and permissions in the past). As we have seen, there are two temporal dimensions to be tackled: the time a norm is valid (when the norm enters the normative system) and the time it is effective (when the norm can produce legal effects). As a consequence, multiple versions of the normative system are needed [30].

The logical machinery used to represent normative change in this approach is complex given that the default logic has to gather very different sorts of default rules providing information on: the content of rules, meta-rules regarding the applicability of other rules, preference between rules, and the timeframe of applicability. For instance, an “abrogation” of a default rule is represented by the addition of a defeater, which is a default rule of a higher order with void content, that is, from which no obligation or permission is derived.

For the example on the regulation of abortion above, the syntactic approach could be roughly illustrated by indicating that in the case of sexual abuse, rule (22) is *not applicable*, and therefore rule (23) is not derived. This could be achieved by introducing a sort of meta-rule to the normative set stating:

24. In the case of sexual abuse, rule (22) is not applicable.

Such a rule would be a *defeater* because it would block the derivation of consequences from rule (22) without excluding it from the normative system. Notice that it adds no normative content by itself.

It is also possible to strengthen the contention that abortion is permitted in the case of sexual abuse by adding another rule to the normative system stating:

25. Abortion is permitted if the pregnancy is the result of sexual abuse.

In Governatori and Rotolo’s approach, this addition is obtained by turning a defeater into a default rule that blocks the application of the original prohibition, but also derives the content of a permission in the case of sexual abuse.

This representation, however, does not capture the basic intuition that derogation is a sort of exclusion where the rule ceases to be a part of the normative system. Instead, since the model concerns the dimension of the efficacy of obligations, a derogation is captured only by blocking the effects of a default rule. Besides, what can be derived depends on which rules are valid at the time when we do the derivation. Thus, in order to keep track of norm changes, Governatori and Rotolo represent different versions of a legal system.

In order to reduce such complexity, Governatori *et al.* [31] explored three AGM-like [6, 7] contraction operators to remove rules, add exceptions and revise rule priorities. Governatori *et al.* [29] also explored a model where, on particular occasions, normative change is reduced to a change of preference relations between default rules.

To illustrate this second method, which may be called the *preferential approach*, consider that from a moral order or a set of constitutional values one may derive inconsistent standards regarding abortion. One may derive permission of abortion from moral considerations, or from arguments about constitutional values, regarding the axiological contention that “women are free to dispose of their own bodies”. But one may also derive prohibition of abortion (rule 22) from a moral contention, or from a constitutional value, stating that “all human beings are the subject of moral worth” and the determination that a “foetus is a human being”.

Hence, this normative system would include rule (22) as well as the following rule:

26. Abortion is permitted.

The presence of rules (22) and (26) makes the normative system inconsistent, and thus the determination of the consequences of these conflicting rules for each relevant circumstance would depend on the addition and change of preference rules such as:

27. In the case of sexual abuse, rule (26) is preferred over rule (22).

In these two alternatives for representing change (syntactic and preferential), the corresponding logic cannot be classical (in particular, it cannot be monotonic). Otherwise rule (22) would conflict with rule (25) and rule (26), thereby making the normative system trivial. In these descriptions, rules (22), (25) and (26) are part of the normative system as “defaults”, and there may be circumstances where each of these becomes inapplicable, or where one of them prevails over another. With the syntactic approach, normative change is a matter of adding new defaults or defeaters to block or enable the normative effects of the defaults over time and according to relevant factors or circumstances. With the preferential approach, normative change is reduced to changing the preference relations between default rules on particular occasions.

In both the syntactic and preferential approaches, a change in the normative system should include not only information about the content of the rules that are subject to change but also information about the applicability of these rules. It is this information about applicability and preference that determines the set of

obligations and permissions derivable from the normative system. Actually, in both these approaches, the set of obligations and permissions may change without any modification to the content of the rules belonging to the normative system. It may be the result of modification to the time span of the applicability of the rules in that set, or the result of a change in the preference relations between defaults.

A third approach, which may be called the *AGM approach*, represents derogation and enactment, respectively, as effective exclusions and additions of content to the normative system. Historically, this was the first approach to modelling normative change, and was originally proposed by Alchourrón and Makinson [6, 7]. When Gärdenfors joined (at that time he was mainly working on counterfactuals), the trio became the founders of the well-known AGM theory, and started the fruitful research area of belief revision [5], which has found many applications in computer science and epistemology. Belief revision is the formal study of how a theory (a deductively closed set of propositional formulas) may change in view of new information that may cause inconsistency with existing beliefs. The basic operations of belief change are expansion (which corresponds to the promulgation of a rule to a code), revision (which corresponds to amendment of the code) and contraction (which corresponds to derogation of its normative application).

One of the first attempts to specify the AGM framework to tackling normative change was put forward by Maranhão [46, 47]. Maranhão introduced a *refinement* operator, which restricts the acceptance of new input to certain conditions in a revision, or keeps a more refined (weaker) version of a rule to be excluded in a contraction. Refinement thus represents the introduction of exceptions to rules in order to avoid conflicts in normative systems (see section 3.6).

More recently, Boella *et al.* [16] also reconsidered the original inspiration for the AGM theory of belief revision as a framework for evaluating the dynamics of rule-based systems. They observed that if we wish to weaken a rule-based system from which we derive too much, we can use the theory of belief base dynamics [34] to select a subset of the rules as a contraction of the rule-based system. Base contraction seems to be the most straightforward and safe way to perform a contraction; it always results in a subset of the original base. But it sometimes means removing too much. In turn, AGM theory contraction may retain some implications of the rule to be deleted. This was one of the motivations for the present contribution. Another advancement is to represent normative change in a formal framework that clearly distinguishes between the concepts of the pertinence of a rule in a normative system and the effectiveness of an obligation in a given context using the input/output logic framework developed by Makinson and van der Torre [42]. A similar approach was proposed by Stolpe [73]. In that work, AGM contractions and revision are used to define derogation and amendment of norms. In particular, the derogation operation

is an AGM partial meet contraction obtained by defining a selection function for a set of norms in input/output logic. Norm revision defined via the Levi Identity characterises the amendment of norms. Stolpe can thus show that derogation and amendment operators are in one-to-one correspondence with the Harper and Levi Identities as inverse bijective maps (cf. section 2.1). Also, Tamargo *et al.* [74, 75] recently studied AGM-like revision operators that consider rules indexed by time intervals.

In the AGM approach, the operation of normative change is performed on the normative system (the set of rules that may be closed under logical consequence). The rules in the original system or in the system resulting from change does not carry meta-information about their applicability, time span or hierarchy (although these features may be added). Therefore, the set of applicable obligations or permissions at a given moment in time is the set of all logical consequences of the normative system valid at that specific time. Hence, information about hierarchy and the time span of validity and applicability is not part of the representation of its rules and does not interfere with the derivation rules of the underlying logic (although such information might be relevant to the revision functions).

To illustrate the AGM approach to the example of abortion discussed above, the normative change would consist in refining rule (22) with respect to the defeating factor “pregnancy resulting from sexual abuse”, resulting in a normative system where rules (23) and consequently (22) are deleted and containing the following rules:

- 25. Abortion is permitted if the pregnancy is the result of sexual abuse.
- 28. Abortion is forbidden if it is not the case that the pregnancy is the result of sexual abuse.

With this last approach, every normative change, that is, every change in the set of obligations and permissions derived from the normative system, amounts to a change to the content of the rules that belong to the set of norms. This aspect makes the set of obligations and conditions for their application closer to the content of the revised normative system.

Research on formal models of normative change has also been concerned with representing legal interpretation.

In the field of artificial intelligence & law, legal interpretation has been mainly explored with models of case-based reasoning, where teleological reasoning is represented to derive solutions to new cases based on precedents. Following Berman and Hafner [13], AI & Law research on teleological reasoning has provided multiple models of the relationship between cases, the factors that such cases include or express,

and the values at stake. Bench-Capon and Sartor [12] assign values to factors, and consequently to rules embedding such factors, to explain precedents according to the applicable rules and the importance of the values promoted by such rules. Prakken *et al.* [61] formalise teleological reasoning using logics for defeasible argumentation, extended to allow the possibility of expressing arguments about values, supported by cases. Sartor [69] explores the proportional balance of constitutional rights, where a legal outcome is compared to alternative outcomes based on their impact on the promotion and demotion of values. He examines the level of consistency between value-based decisions of cases given the factors present in those cases [70].

In turn, AI & Law research on statutory interpretation has focused on the dynamic ascription of meanings to rules. These contributions are based on the distinction between “*constitutive*” (or “*conceptual*”) rules ascribing meanings to facts or objects and “*regulative*” rules demanding, prohibiting or permitting actions or states [32]. Interpretation is then modelled as introducing or changing conceptual rules. Governatori and Rotolo [30] represent such changes, within the syntactic approach, as the introduction of exceptions, by blocking the application of default rules to a given condition or constellation of factors. Boella *et al.* [15] developed that model by introducing values as coherence parameters guiding the change of conceptual rules, parameters whose meanings may be extended (weakening the antecedent of a conditional rule) or restricted (strengthening the antecedent of a conditional rule).

The incorporation of the AGM approach into input/output logics [16] and, later, the representation of normative systems in an architecture of input/output logics combining constitutive and regulative rules, brought a new perspective to representing legal interpretation [18]. Maranhão and de Souza [52] introduced a contraction function for such combined normative sets in order to represent choices in legal doctrine between changing the definitions (or meaning ascriptions) of legal terms and changing the content of legal regulative rules, taking into consideration the network effects of those changes.

Maranhão [50] proposed an architecture of input/output logics for modelling doctrinal interpretation where values are represented as rules, and constitutive and regulative rules are the object of different contraction, revision and refinement functions. Differently from the work of Boella *et al.* [15], where legal interpretation is conceived as a dynamic of syntactic modifications to constitutive rules (within the syntactic approach), in Maranhão’s model it is not only constitutive rules, but also values and regulative rules, that are subject to change (with the AGM approach) in order to reach a coherent and stable description of the normative system. More recently, Maranhão and Sartor’s [53] research on statutory interpretation built on the case-based tradition of teleological reasoning and balancing with their repre-

sensation of legal construction—where a model of balancing values is incorporated into an architecture of input/output logics—serving as a reference to the revision of constitutive (meaning ascriptions) and regulative rules.

Which is the best approach to representing normative change—syntactic, preferential or AGM?

This question was controversial in the 1990s in the context of Alchourrón's [3] criticism that defeasible logics are philosophically inadequate. According to Alchourrón, defeasible logic unnecessarily weakens the inferential power of the underlying logic. It obscures the fact that the defeat of a conclusion is actually the result of the dynamic of revising the premises in a derivation, or the fact that the defeat of a consequence results from revising the antecedent of a conditional. According to Alchourrón, in an adequate account of the epistemology of law or of any domain, the revision processes of the premises of an argument or the antecedent of a conditional should be transparent [48].

Actually the reply to this question depends on what aspect of legal reasoning one would like to capture with the model of representation (without considering the technical issue of computational complexity).

As we have seen, there is a fundamental difference between the pertinence of a rule to a normative system and its effects in terms of the derivability of the corresponding obligations/permissions in the presence of given circumstances. There is the time span for when a rule pertains to the normative system, that is, the time the rule exists in the normative system. But, although pertinent to a system, a rule may still not produce its effects, for example because its conditions of application are dependent on an event or regulated by another rule, so there is another time span for when the norm is applicable. Furthermore, as mentioned above, there is the time span for when the conclusions of an instantiated rule apply to a particular individual, considering that the instantiated rule may be derogated or annulled (i.e. declared invalid) for that particular individual by a judicial authority.

The distinction between the validity and efficacy of a rule may be captured by all approaches. But the syntactic approach seems to be more congenial to the dimension of efficacy, that is, the applicability of rules, considering that the revision operations are represented as syntactical changes to the rules that affect their applicability. A contraction operator does not properly exclude a rule but interferes with the derivability of its consequence.

In turn, the AGM approach seems to be more congenial to modelling the dynamics of the pertinence of a rule in a normative system, since the suppression or addition of obligations or permissions, and obligations derived from the basic set of rules, are reflected in proper exclusion or expansion to the rules of the normative system.

In the end, the description of the obligations and permissions derived from the normative system may coincide in both approaches, the difference lying in the set of basic rules.

Lastly, the preferential approach seems to be more congenial to the dynamic of legal principles and values related to positively enacted rules. Such principles and values, both considered as external to the normative system or enshrined in the constitution, potentially conflict but coexist in the normative order or political morality underlying such an order of legal rules. Depending on the context, they are balanced in order to derive a solution. The preferential approach reflects the fact that the derivation of a normative solution from principles or values results from resolving potential conflicts by giving more weight to a preferred principle than another principle in a given context.

It seems that a closer correspondence between the content of the rules and the applicable obligations/permissions is also of interest for the representation of legal construction where a particular reconstruction of a fragment of the normative system takes place before the instantiation of an operational rule.

Recent research on models of legal interpretation has shown that the three approaches must be combined since, as we have seen, the interpretive activity, particularly legal construction, involves all of the following three dimensions:

- manipulation and refinement of constitutive and regulative rules in a normative system (*validity*);
- consideration and weighing of underlying values (*balancing*);
- adaptation of definitions of legal terms to make the rules isomorphic and applicable to the facts of a particular case (*applicability*).

The first two approaches listed in this section are presented in the work of Tamargo *et al.* [75]. This article focuses on the AGM option, presenting its reformulation for input/output logics—a family of logics dedicated to the analysis of normative reasoning in particular as well as rule-based reasoning in general. We consider the combination of these two formal approaches, AGM belief change and input/output logics, to be a promising framework for analysing normative change. On the one hand, the kind of analysis of information change that AGM-like approaches pursue is insightful and very clear at the same time, and often can be reformulated into specific solutions for other formal frameworks. On the other hand, input/output logics offer an analysis of rule-based reasoning that is along the same lines, since it combines the immediate clarity of characterising distinct rule-based systems via the structural properties they satisfy with an in-depth analysis of the different

kinds of rule-based reasoning that can be modelled. In our view, applying an AGM-like approach on top of input/output systems allows an essential characterisation of change to be developed that focuses here on normative reasoning, but can actually be extended to other forms of rule-based reasoning.

2 Formal Framework

In this section, we briefly introduce the formal framework we will adopt in our analysis of normative change. In the last few decades, the area of knowledge representation and reasoning has proposed various formal approaches to modelling the dynamics of knowledge, and to modelling normative change in particular. As a result, one methodological issue that we need to address is what kind of analysis do we want to develop for normative change.

2.1 The AGM Approach

We will rely on the methodology of the *AGM approach* to belief change that we introduced in section 1.4. In the last 30 years, AGM has been the most popular formal approach to analysing the dynamics of beliefs, but it has been debated whether it is the best approach to analysing belief change in general, and normative change in particular. In this section, we briefly outline the main characteristics of this approach for the unfamiliar reader, and discuss why we still consider it to be a viable option for analysing normative change.

Let's start with a well-known example. Our knowledge base contains the following information:

- a.* Sweden is an European country.
- b.* All European swans are white.
- c.* The bird I just caught in the trap is a swan.
- d.* The bird I just caught in the trap is from Sweden.
- e.* No bird can be black and white at the same time.

This information entails that the bird I just caught in a trap is white. But then I look at it and I see that it is undoubtedly black. I add to my knowledge base the following proposition:

- f.* The bird I just caught in the trap is black.

From my knowledge base, I must conclude that the bird I just caught in the trap is both white and black. My knowledge base contains conflicting information, it is inconsistent. How should the situation be fixed? What constraints should we follow in changing our beliefs? And how should we give a formal characterisation to such constraints?

It is generally assumed that the constraints that a *rational* form of belief change should respect are based on considerations of two kinds:

1. *Logic*. Here the focus is on *consistency preservation*: the content of our knowledge base should always be devoid of contradictions.

Looking at our example, we cannot accept that we can believe that a bird is black and that the bird is white at the same time. Once we rely on piece of information f , we need to change the content of our knowledge base, since propositions a - f together necessarily imply a contradiction.

2. *Pragmatic*. This point and Point 1 above are intertwined. If we are forced to modify the content of our knowledge base in order to satisfy logical constraints, e.g. in order to preserve consistency, we should do so taking into consideration also pragmatic issues, based on, for example, *economy of information*. According to that principle, information is valuable, some pieces of information are more relevant and reliable than others, and if we are forced to drop some pieces of information, we should “minimise the damage” by eliminating only the minimal amount of information that is necessary to preserve logical consistency.

What should we do in our example once we learn proposition f and we spot the conflict? We could simply erase the entire knowledge base, just eliminate all the propositions (a)-(e). But why should we do this given that, for example, it is sufficient to drop only one proposition among (a), (b), (c), (d), and (e)?

In order to describe belief change, the AGM approach gives a formal definition to the knowledge representation desiderata by defining formal constraints based on logical or pragmatic considerations.

To formally introduce the AGM approach, we need some formal preliminaries. We use a classical propositional language $\mathcal{L}$, built from atomic propositional letters and using the propositional connectives $\neg, \wedge, \vee, \rightarrow, \equiv, \perp$. Lower-case letters $a, b, c, \dots, x, y, z$ will be used to represent propositions. A *knowledge base* is a set of propositional formulas, that will be indicated by capital letters as $\mathcal{K}$. In addition, $\models$ and Cn will represent the classical propositional entailment relation and entailment operator respectively.

The epistemic status of an agent is characterised by a knowledge base $\mathcal{K}$. Actually, the classical AGM approach embraces a perspective that has been dominant in epistemic logics: the epistemic status of the agent is characterised using a *belief set*, a logical theory closed under Cn . That is, the epistemic status of an agent is characterised by a knowledge base $\mathcal{K}$ such that $\mathcal{K} = Cn(\mathcal{K})$. Let $\mathcal{T}$ be the set of the belief sets (i.e. the closed theories) of language $\mathcal{L}$, that is, $\mathcal{T} := \{\mathcal{K} \subseteq 2^{\mathcal{L}} \mid \mathcal{K} = Cn(\mathcal{K})\}$.

The first question we need to address is what kind of changes we should consider. The AGM approach recognises three operations as the basic ones: *expansion*, *contraction*, and *revision*. Assume our agent A has a knowledge base $\mathcal{K}$:

- *Expansion* $+$: A is informed that proposition p holds, and simply adds it to $\mathcal{K}$ without caring whether this could generate some contradiction. The resulting knowledge base is indicated as $\mathcal{K} + p$.
- *Contraction* $-$: A believes that p holds ($p \in \mathcal{K}$), but then decides that it is better to abandon such a belief, for example because the source is not considered trustworthy anymore. The resulting knowledge base, indicated as $\mathcal{K} - p$, should be such that p is no longer implied by A 's knowledge base.
- *Revision* $*$: A is informed that proposition p holds, and wants to add it to $\mathcal{K}$, but with the proviso that the resulting knowledge base should be logically sound. The resulting knowledge base is indicated as $\mathcal{K} * p$.

These three kinds of operations can be characterised using the function

$$\bullet : \mathcal{T} \times \mathcal{L} \mapsto \mathcal{T} \text{ with } \bullet \in \{+, -, *\}.$$

Actually, the truly basic operations are generally considered to be the first two, *expansion* and *contraction*, since *revision* is usually built on top of those using the so-called *Levi Identity* [40]:

$$\mathcal{K} * p := (\mathcal{K} - \neg p) + p.$$

Revising knowledge base $\mathcal{K}$ by introducing a new proposition p requires that we guarantee that there are no pieces of information in our knowledge base that are in conflict with p . The reasonable way of obtaining this is to contract $\mathcal{K}$ to ensure that it does not imply $\neg p$, and only then introduce p . This is the revision procedure that is modelled by the Levi Identity.

In the swan example, in order to revise the belief set with the information that the swan is black, we should proceed as follows: the belief set corresponds to the set $\mathcal{K} := Cn(\{a, b, c, d, e\})$ and we want to introduce f ("The swan in the trap is black"). Using the Levi Identity, the revision

$$\mathcal{K} * f$$

will consist in first contracting the piece of information $\neg f$ (“It is not the case that the swan in the trap is black”) from $\mathcal{K}$. The resulting belief set, $\mathcal{K} - \neg f$, should be a set of formulas that is smaller than $\mathcal{K}$ and does not imply $\neg f$ anymore. For example, let us opt for weakening proposition b (“All European swans are white”) into a new proposition b' (“All European swans are white or black”), that is, $\mathcal{K} - \neg f = Cn(\{a, b', c, d, e\})$, and it is easy to check that $\mathcal{K} - \neg f$ does not imply $\neg f$ anymore. Only after the contraction do we add f , that is, we can set $\mathcal{K} * f = (\mathcal{K} - \neg f) + f = Cn(Cn(\{a, b', c, d, e\}) \cup \{f\})$, that is, $\mathcal{K} * f = Cn(\{a, b', c, d, e, f\})$.

We also have a complementary construction, the *Harper Identity*, in which revision is the primitive operator and contraction is defined on top of it:

$$\mathcal{K} - p := (\mathcal{K} * \neg p) \cap \mathcal{K}.$$

$\mathcal{K} - p$ should be a subset of $\mathcal{K}$ not implying p , while $\mathcal{K} * \neg p$ should be a theory as close as possible to $\mathcal{K}$ that implies $\neg p$ and does not imply p . The meaning of the Harper Identity is that since $\mathcal{K} * \neg p$ should not imply p , if we intersect it with $\mathcal{K}$, we obtain a contraction: a subset of $\mathcal{K}$ that does not imply p .

We can rephrase the above example to show that the Harper Identity and the Levi Identity can correspond to each other. Let $\mathcal{K}$ be our knowledge base containing propositions (a)-(e), and assume that we have a revision operator $*$, as described above and which is introduced here as a primary operator, such that $\mathcal{K} * f = Cn(\{a, b', c, d, e, f\})$. If we use the Harper Identity to define a contraction operator $-$ from $*$, we obtain $\mathcal{K} - f = Cn(\{a, b', c, d, e, f\}) \cap Cn(\{a, b, c, d, e\})$ that, since $b \models b'$, corresponds to $\mathcal{K} - f = Cn(\{a, b', c, d, e\})$, that is, the contraction we have used above as a primitive operator to define $*$ via the Levi Identity. In what follows, we will use both Levi and Harper Identities, and we will soon give a more formal definition of the correspondence between the two.

Once we have identified the basic operations we are interested in, the second question we need to address is how we want to model and constrain such change operations. For each kind of operation, we want to determine a set of desired properties they should satisfy, and give a formal expression to such desiderata.

Expansion is considered to be a trivial operation, formalised by adding the formula we are interested in to the knowledge base and letting the agent commit to all the logical consequences of such an addition:

$$\mathcal{K} + a := Cn(\mathcal{K} \cup \{a\}).$$

In the *contraction* operation, an agent starts with a belief set $\mathcal{K}$ (e.g. the theory determined by sentences (a)-(e) above) and wants to eliminate some pieces of information in the belief set (e.g. that the swan is white). The AGM approach gives a formal representation to a basic set of desiderata using six *postulates*.

Definition 6 (AGM contraction [5]). *Let $-$ be a function that, given a belief set $\mathcal{K}$ and a proposition a , returns a new belief set $\mathcal{K} - a$. Function $-$ is an AGM basic contraction operator iff it satisfies the following postulates:*

- (- 1) $\mathcal{K} - a$ is closed under Cn (closure)
- (- 2) $\mathcal{K} - a \subseteq \mathcal{K}$ (inclusion)
- (- 3) If $a \notin \mathcal{K}$, then $\mathcal{K} - a = \mathcal{K}$ (vacuity)
- (- 4) If $\not\models a$, then $a \notin \mathcal{K} - a$ (success)
- (- 5) If $a \in \mathcal{K}$, then $\mathcal{K} \subseteq (\mathcal{K} - a) + a$ (recovery)
- (- 6) If $\models a \equiv b$, then $\mathcal{K} - a = \mathcal{K} - b$ (extensionality)

Two extra postulates are introduced to relate the contraction of complex formulas to the contraction of their components:

- (- 7) $\mathcal{K} - a \cap \mathcal{K} - b \subseteq \mathcal{K} - (a \wedge b)$ (conjunctive overlap)
- (- 8) If $a \notin \mathcal{K} - (a \wedge b)$, then $\mathcal{K} - (a \wedge b) \subseteq \mathcal{K} - a$ (conjunctive inclusion)

Function $-$ is an AGM contraction operator iff it satisfies postulates (- 1)-(- 8).

We will briefly go through the meaning of these postulates. Postulate (- 1) enforces an idealisation we have already discussed: the epistemic status of the agent is described using logically closed theories (belief sets), hence every change operation must transform a closed theory into a new closed theory. Postulate (- 2) imposes that the change operation must result in an actual *contraction* of the agent's belief set, that is, the set of formulas believed by the agent at the end is a subset of the initial beliefs. Postulate (- 3) formalises a principle of an economical nature: if the contraction operation involves a formula that is already excluded from the agent's beliefs, the contraction operation is *vacuous*, that is, nothing changes, since the desired result is already satisfied. Postulate (- 4) imposes that, whenever possible, that is, whenever the formula to be contracted is a *contingent* formula and not a

tautology, the contraction operation must be successful, that is, the formula should no longer be in the resulting belief set. Let us jump to postulate (– 6), leaving postulate (– 5) aside for one moment. Postulate (– 6) imposes independence from syntax, which is a classical logical principle: whenever two pieces of information are logically equivalent, they are indifferent from a logical point of view, and their impact on the agent’s belief set is exactly the same. It is easy to see that this principle is strongly related to postulate (– 1), the use of logically closed theories to model the epistemic states. While the use of closed theories imposes indifference with regard to the syntactic form of the knowledge base in the *static* model of the agent’s epistemic state, the principle of *extensionality* extends such syntactic indifference also to operations modelling the *dynamics* of the agent’s epistemic states. Postulates (– 7) and (– 8) are considered extra postulates, since they are the only ones that impose constraints on the way a contraction operator behaves with different formulas, in particular how the contraction of a formula should behave with the contraction of logically weaker formulas.

Postulate (– 5), *recovery*, has a special status, since, probably together with postulate (– 1), it is the most debated AGM principle, and in a certain sense it is also the one that mainly characterises the classical AGM approach. Its nature is purely economical, based on the idea that in order to contract, we “cut” as little as possible from the original knowledge base. So little that if the agent decides that contracting by formula a was not a good idea and that a should be added back, we should be able to return to the original knowledge base without any loss. In fact, together with postulate (– 2), postulate (– 5) implies that if $a \in \mathcal{K}$, then $\mathcal{K} = (\mathcal{K} - a) + a$, that is, if we put a back after a contraction, we go back to the initial state. It has been debated extensively whether *recovery* is a reasonable principle for contraction, and we will return to this issue later in this section.

Anyway, the reader can see that each of these eight postulates answers to either logical or pragmatic desiderata. For a more detailed explanation of their meaning, we refer the interested reader to the original AGM paper [5] and many other publications in the field.

It is worth mentioning that Rott [67] has disputed whether the AGM approach does actually satisfy any principle of informational economy. Despite the relevance of Rott’s observations, postulates like (– 3) and (– 5) are generally seen as necessary conditions for defining contraction operators that satisfy the principle of informational economy. The principle of informational economy, which has been expressed in various forms and with different names, has always been addressed by researchers in the area as the main guideline for the definition of postulates.

In our presentation of AGM belief change, we first introduced a set of possible change operations (specifically, *expansion*, *contraction*, and *revision*), and then a set

of *postulates* to give formal expression to the properties we think such operations should satisfy, specifically those for contraction. The next step is to present the formal tools that we can use to define such change operators. That is, given a set of postulates, the AGM approach is focused on providing a formal characterisation of the class of operations that satisfy such postulates. The classical results in the area define classes of change operations using maxiconsistent subsets and choice functions [5], orderings over possible-world semantics representing which situations the agent considers to be more plausible [33, 37], or orderings over the formulas (*epistemic entrenchment relations*) indicating which pieces of information the agent considers to be more or less reliable [27].

Regarding contraction, the initial characterisation of the class of operations satisfying the basic postulates is based on identifying the maximal subsets of the belief set that do not imply the contracted formula. The resulting belief set is defined by the intersection of some such maximal subsets. Which maximal subsets are used in the definition of the contraction is formalised via a dedicated choice function.

Definition 7 (Partial meet contraction [5, p. 512]).

Let $\mathcal{K} \perp a$ be the remainder set, containing the maximal subsets $\mathcal{K}'$ of $\mathcal{K}$ such that $\mathcal{K}'$ is a closed theory and $a \notin \mathcal{K}'$. That is, $\mathcal{K}' \in \mathcal{K} \perp a$ iff

- (i) $\mathcal{K}' \subseteq \mathcal{K}$,
- (ii) $\mathcal{K}' \in \mathcal{T}$,
- (iii) $a \notin \mathcal{K}'$, and
- (iv) there is no set $\mathcal{K}'' \in \mathcal{T}$ such that $\mathcal{K}' \subset \mathcal{K}'' \subseteq \mathcal{K}$ and $a \notin \mathcal{K}''$.

Let pm be a choice function defined over the set of the remainder sets. Function pm is a partial meet function if for every KB $\mathcal{K}$ and every formula a :

- $pm(\mathcal{K} \perp a) \subseteq \mathcal{K} \perp a$, and
- if $\mathcal{K} \perp a \neq \emptyset$, then $pm(\mathcal{K} \perp a) \neq \emptyset$.

A partial meet contraction operator $-$ is defined as: $\mathcal{K}_A^- = \bigcap pm(\mathcal{K} \perp A)$.

The class of partial meet contractions is sufficient to give an operational characterisation of the class of AGM basic contraction operations.

Observation 8. [5, Observation 2.5] A contraction operator $- : \mathcal{T} \times \mathcal{L} \mapsto \mathcal{T}$ is an AGM basic contraction operator (satisfying (-1) – (-6)) iff it is a partial meet contraction operator.

An analogous analysis can be developed for *revision*. First of all, we can formalise our desiderata via appropriate postulates.

Definition 9 (AGM revision $*$ [5]). *Let $*$ be a function that, given a belief set $\mathcal{K}$ and a proposition a , returns a new belief set $\mathcal{K} * a$. Function $*$ is an AGM basic revision operator iff it satisfies the following postulates:*

- ($*$ 1) $\mathcal{K} * a$ is closed under Cn (closure)
- ($*$ 2) $a \in \mathcal{K} * a$ (success)
- ($*$ 3) $\mathcal{K} * a \subseteq \mathcal{K} + a$ (inclusion)
- ($*$ 4) If $\neg a \notin \mathcal{K}$, then $\mathcal{K} + a = \mathcal{K} * a$ (vacuity)
- ($*$ 5) $\perp \in (\mathcal{K} * a)$ iff $\models \neg a$ (triviality)
- ($*$ 6) If $\models a \equiv b$, then $\mathcal{K} * a = \mathcal{K} * b$ (extensionality)

Two extra postulates are introduced also for revision. These postulates relate the revision of complex formulas to the revision of their components:

- ($*$ 7) $\mathcal{K} * (a \wedge b) \subseteq (\mathcal{K} * a) + b$ (Iterated ($*$ 3))
- ($*$ 8) If $\neg b \notin \mathcal{K} * (a)$ then $(\mathcal{K} * a) + b \subseteq \mathcal{K} * (a \wedge b)$ (Iterated ($*$ 4))

Function $*$ is an AGM revision operator iff it satisfies the postulates ($*$ 1)-($*$ 8).

The meaning of the postulates for revision is very close to the meaning of the postulates for contraction. The parallel is clear for postulates ($*$ 1), ($*$ 2), ($*$ 3), ($*$ 4), ($*$ 6) and the correspondent postulates for contraction. Postulate ($*$ 5) imposes perhaps the key rational desideratum for modelling belief dynamics: preserving consistency. Whenever we add a new piece of information a , the only case where the resulting belief set can be inconsistent is when a itself is inconsistent.

We briefly summarise a series of well-known basic results in the area that show how the notions introduced up to this point are solidly connected to one other in AGM theory. First of all, the construction of AGM revision and contraction operators are intertwined via the Levi Identity.

Observation 10. [5] *Let $*$: $\mathcal{T} \times \mathcal{L} \mapsto \mathcal{T}$ be a revision operator. Function $*$ is a basic AGM revision operator (it satisfies ($*$ 1)-($*$ 6)) if and only if there is a contraction operator $-$ such that:*

- $*$ can be defined via the Levi Identity from $-$. That is, for every $\mathcal{K}$ and a ,

$$\mathcal{K} * a = (\mathcal{K} - \neg a) + a$$

- $-$ is a basic AGM contraction operator (it satisfies (-1) – (-6)).

Given Observation 8, Observation 10 connects the construction of basic AGM revision operators to the class of partial meet contractions via the Levi Identity.

An analogous result [5] holds for contraction and revision operators satisfying postulates (-1) – (-8) and $(*1)$ – $(*8)$ respectively.

Such a dependency of revision on contraction can also be reversed, moving from AGM revision operators to the definition of AGM contraction operators: the one-to-one correspondence between the Levi Identity and the Harper Identity, that we have briefly exemplified above in revising and contracting our knowledge base about swans, can actually be formally proved. Let us translate the Levi and Harper Identities into transformation functions. Given a belief set $\mathcal{K}$, a formula a , a contraction operator $-$ and a revision operator $*$, let

- $\mathcal{K} \mathbb{R}(-) a := Cn((\mathcal{K} - \neg a) \cup \{a\})$
- $\mathcal{K} \mathbb{C}(*) a := (\mathcal{K} * \neg a) \cap \mathcal{K}$

where $\mathbb{R}(-)$ represents a revision operator obtained from contraction $-$ via the Levi Identity and $\mathbb{C}(*)$ represents a contraction operator obtained from revision $*$ via the Harper Identity. Using these operators, Makinson has proven that there is full correspondence between the Levi and Harper Identities.

Observation 11. [41] *Let $\mathcal{K}$ be a belief set, and let a be a formula, with $\mathbb{R}(-)$ and $\mathbb{C}(*)$ defined as above.*

- *Let $-$ satisfy the postulates of closure, inclusion, vacuity, extensionality, and recovery. Then $\mathbb{C}(\mathbb{R}(-)) = -$.*
- *Let $*$ satisfy the postulates of closure, inclusion, success, and extensionality. Then $\mathbb{R}(\mathbb{C}(*)) = *$.*

As an immediate consequence, the Levi and Harper Identities have been shown to be interchangeable for AGM theory:

$$\mathcal{K} * a = (K \cap K * a) + a;$$

$$\mathcal{K} - a = K \cap ((\mathcal{K} - a) + \neg a).$$

What we have presented up to this point are some key results of the AGM approach that provide an essential introduction to the unfamiliar reader, and which are relevant to the sections that follow.

2.2 Criticisms of the AGM Approach

Simplifying, we could say that there are three main steps that characterise the AGM method:

- the identification of the typologies of change we want to model and of the properties we want them to satisfy;
- the translation of such desiderata into postulates, that is, into formal constraints;
- the characterisation of the classes of operators that satisfy the desired set of postulates. Such a characterisation is usually obtained by proving the correspondence of such operators to a class of constructions defined using a relevant formal tool (e.g. maxiconsistent sets, possible-world models...).

The AGM approach to belief change has quickly become standard in the field, and the last 30 years has seen many contributions [25]. Despite the fact that it has become a major research topic in knowledge representation, it is an approach that has been frequently and heavily criticised, and new lines of research have sprouted from some of these critiques. We briefly list some of the main critiques the AGM approach has received.

2.2.1 Too Many Constraints Imposed on the Underlying Logic

The AGM approach was originally developed for classical propositional logic (PL), and the classical results assume that the underlying logic, characterised by a language L and an entailment operator Cn , satisfies many of the formal properties that characterise PL:

1. The language L is closed under the propositional connectives.
2. The entailment operator Cn is *Tarskian*, that is, given two sets of formulas $\mathcal{K}, \mathcal{K}' \subseteq L$, it satisfies the following properties:
 - *monotonicity*: if $\mathcal{K} \subseteq \mathcal{K}'$, then $Cn(\mathcal{K}) \subseteq Cn(\mathcal{K}')$;
 - *idempotence*: $Cn(\mathcal{K}) = Cn(Cn(\mathcal{K}))$;
 - *iteration*: $\mathcal{K} \subseteq Cn(\mathcal{K})$.
3. The consequence operator satisfies some well-known properties of classical logic:

- *deduction*: $b \in Cn(\mathcal{K} \cup \{a\})$ iff $(a \rightarrow b) \in Cn(\mathcal{K})$;
- *disjunction in the premises*: if $a \in Cn(\mathcal{K} \cup \{b\})$ and $a \in Cn(\mathcal{K} \cup \{c\})$, then $a \in Cn(\mathcal{K} \cup \{b \vee c\})$.

4. *Compactness*: if $a \in Cn(\mathcal{K})$, then $a \in Cn(\mathcal{K}')$ for some finite $\mathcal{K}' \subseteq \mathcal{K}$.

Much recent research in belief revision has been dedicated to investigating whether the above constraints are essential to the definition of AGM operators and, when we are dealing with an underlying logic that does not allow the definition of classical AGM postulates, what other meaningful postulates can be defined and satisfied. For example, the AGM approach has been applied to logics that are not fully closed under propositional operators [21, 80], that are not monotonic [79, 20, 19], and that are not compact [65].

This article will also deal with a family of logics that do not satisfy all the properties listed above. Input/output logics are not closed under propositional operators and, because of that, cannot satisfy properties like *deduction* and *disjunction in the premises*. Some input/output logics also do not satisfy the property of *monotonicity* [43]. Although we shall not discuss them in this article, application of the AGM methodology to normative change based on non-monotonic input/output logics is a promising field of inquiry.

2.2.2 Lack of Expressiveness

It has often been pointed out that the expansion/contraction/revision triad is not sufficient to account for the dynamics of information. It is also claimed that the AGM approach is not appropriate for handling multi-agent systems because it is suitable only for factual information.

With respect to the first line of criticism, it is worth mentioning that operations that are not reducible to the original ones have been introduced, such as *update* [36] and *merging* [39] among others. Besides, many refinements to the original operations have been proposed, based on alternative postulates and formal constructions, which introduce new dimensions to the original operations, such as the trustworthiness of the new information [25, Chapter 8]. Despite being a common place that the AGM operations of contraction and revision are not sufficient to cover all the relevant dynamics of information, it is generally accepted that analysing the operations of contraction and revision is a good starting point for modelling informational change in many contexts. Analysing contraction and revision in different formal contexts allows us to deal with the ideas of minimal change and consistency preservation in

each of those contexts, and minimal change and consistency preservation are the two main stepping stones towards characterising rational informational change.

It is true that multi-agent contexts are not immediately compatible with the AGM approach, since some classical AGM postulates would be counter-intuitive in such a framework.

In the area of Dynamic Epistemic Logic (DEL), it has been pointed out that some sentences, for example those resembling the structure of that used in *Moore's paradox*, are not compatible with the *success* postulate [76]. The DEL framework allows us to model the dynamics of epistemic states in which the agent also models higher-order sentences representing beliefs about its own beliefs and the beliefs of other agents. On the other hand, AGM is easier to understand, and allows a more in-depth analysis of specific kinds of operations. Working first at the AGM level, and later transporting the proposed solutions to other frameworks such as the DEL framework, can be seen as a good research strategy. Also, some domains, like formal ontologies or the domain under consideration in this article, normative bodies, do not usually need to deal with a multi-agent aspect in modelling change.

2.2.3 Logical Closure and the Recovery Postulate

Finally, let us consider two further lines of criticisms of the AGM approach that are particularly relevant for what follows. These are connected to the *recovery* (-5) and the *closure* ($((-1)/(*1))$) postulates.

As mentioned above in this section, the recovery postulate has often been criticised. On the one hand, its desirability is intertwined with the use of logically closed belief sets. On the other hand, as many commentators have pointed out, the recovery postulate is not always desirable even if we are working with closed belief sets (see [25, Sect. 5.1] for an overview).

Moreover, if we define revision on top of contraction via the Levi Identity, it turns out that the recovery postulate is not necessary to characterise the class of the AGM basic revision operators. That is, the representation that results in Observation 10 remains valid if we drop postulate (-5).

Observation 12. [28] *Let $* : \mathcal{T} \times \mathcal{L} \mapsto \mathcal{T}$ be a revision operator. Function $*$ is a basic AGM revision operator (it satisfies $(*1)$ - $(*6)$) if and only if there is a contraction operator $-$ such that:*

*$\cdot *$ can be defined via the Levi Identity from $-$. That is, for every $\mathcal{K}$ and a ,*

$$\mathcal{K} * a = (\mathcal{K} - \neg a) + a.$$

· – satisfies (-1) – (-4) and (-6) .

The criticisms of the recovery postulate, together with the fact that it is not a necessary property in order to characterise well-behaved revision operators, has convinced many researchers to drop such a postulate in many contexts, looking for more significant alternatives [24].

As mentioned above, the AGM approach models change over belief sets, that is, it does not consider arbitrary sets of formulas, but only logically closed theories.

This is a constraint that is in line with the classical modelling approach of epistemic logics, and it is prone to the same kind of criticisms. On the one hand, characterising epistemic states as closed logical theories is seen as the correct way to characterise rational agents, since it allows a description of knowledge that is syntax-independent and that models the commitment a rational agent should have towards all the consequences of what is explicitly stated in a knowledge base. On the other hand, depending on the modelling goals, exactly the same arguments can be considered as drawbacks. If we investigate the belief states and dynamics of agents with bounded rationality, committing to closed logical theories is too strong an idealisation, which in epistemic logics is labelled as *logical omniscience*. Moreover, the syntactic form of the knowledge base can actually play a role in modelling the way the agent manages the information at its own disposal, for example by making explicit how the agent clusters pieces of information together in a single formula. The belief change community has reacted by developing the theory of *base revision*, where the same approach as AGM to investigation is applied to finite knowledge bases rather than logically closed theories [35].

2.3 Base Contraction and Revision

In base revision, the epistemic status of an agent is described using a set of formulas K that is not necessarily logically closed. The basic operation in base revision is Hansson’s *kernel contraction* [35], which is a re-interpretation at the level of finite base of the AGM notion of contraction based on remainder sets.

Hansson’s base contraction is based on the notions of *kernels* and *incision functions* in a way that resembles the roles of the *remainder sets* and the *partial meet functions* in partial meet contraction. Given a knowledge base K and a formula a , the *a-kernels* of K are the minimal subsets of K that have a as a logical consequence. Eliminating some pieces of information from each kernel allows us to avoid deriving a , and such an elimination is made using an *incision function*.

Definition 13 (Kernel set and incision function [35]). *Let $a \in \mathcal{L}$ and $K \subseteq \mathcal{L}$. The set $\text{Kern}_K(a) \subseteq 2^{2^{\mathcal{L}}}$ is the kernel set of K with respect to a if it is defined as follows. $X \in \text{Kern}_K(a)$ if and only if:*

- $X \subseteq K$;
- $a \in \text{Cn}(X)$;
- if $X' \subset X$, then $a \notin \text{Cn}(X')$.

An incision function σ defined over the kernel sets is a choice function such that:

- $\sigma(\text{Kern}_K(a)) \subseteq \bigcup \text{Kern}_K(a)$;
- $\sigma(\text{Kern}_K(a)) \cap X \neq \emptyset$ for all $X \in \text{Kern}_K(a)$.

Once the incision function has specified the information that should be eliminated from K in order to avoid deriving a , we can use it to define a contraction operator on arbitrary sets of formulas.

Definition 14 (Kernel contraction [35]). *Let $a \in \mathcal{L}$ and $K \subseteq \mathcal{L}$. Operator $-_{\sigma} : 2^{\mathcal{L}} \times \mathcal{L} \mapsto 2^{\mathcal{L}}$ is a kernel contraction operator if*

$$K -_{\sigma} a = K \setminus \sigma(\text{Kern}_K(a)).$$

Hansson gives a postulate characterisation of kernel contractions.

Observation 15. [35] *A function $- : 2^{\mathcal{L}} \times \mathcal{L} \mapsto 2^{\mathcal{L}}$ is a kernel contraction if and only if it satisfies the following postulates:*

- ($-_{\sigma}$ 1) $K - a \subseteq K$ (inclusion)
- ($-_{\sigma}$ 2) If $\nexists a$, then $a \notin K - a$ (success)
- ($-_{\sigma}$ 3) If $b \in K \setminus K - a$, then there is a $K' \subset K$ such that $a \notin \text{Cn}(K')$ but $a \in \text{Cn}(K' \cup \{b\})$ (core-retainment)
- ($-_{\sigma}$ 4) If for all subsets K' of K , it holds that $a \in \text{Cn}(K')$ iff $b \in \text{Cn}(K')$, then $K - a = K - b$ (uniformity)

We can also define revision combining contraction and expansion using bases, but now we have two possible ways of combining the two operations [34],

- $\mathcal{K} *__{\sigma} a = (\mathcal{K} -_{\sigma} \neg a) +_{\sigma} a$ (Levi Identity)
- $\mathcal{K} *__{\sigma} a = (\mathcal{K} +_{\sigma} a) -_{\sigma} \neg a$ (Reversed Levi Identity)

where $\mathcal{K} +_{\sigma} a := \mathcal{K} \cup \{a\}$. The two options define revision operators with different properties [34]. The Reversed Levi Identity is not a viable option when we are working with belief sets, since the first step, the expansion, could take us to an inconsistent theory, the contraction of which is not efficiently managed by the classical AGM approach.

3 Formal Analysis of Normative Change

The distinction between norms and obligations was articulated and formally developed in input/output logic [42]. Input/output logic takes a very general view of the process used to obtain conclusions (outputs) from given sets of premises (inputs). To detach an obligation from a norm, there must be a context, and the norm must be conditional. Thus, norms are just particular kinds of rules, and one may view a normative system simply as a set of rules.

Makinson's iterative approach to normative reasoning distinguishes unconstrained from constrained output. Unconstrained is close to classical logic, whereas constrained output is much less similar, due to the existence of multiple output sets (or extensions), for example. Examples of constrained output are default reasoning, defeasible deontic reasoning etc.

Makinson and van der Torre introduced seven distinct input/output logics, including both a semantic definition and a proof theoretic characterisation [43, 44]. They showed that their seven unconstrained input/output logics cannot handle contrary-to-duty reasoning and thus cannot be used as logics representing normative reasoning. They therefore introduced constrained output in a companion paper, and they showed how that can be used as a logic of norms. However, the user has to make some seemingly arbitrary choices by, for example, choosing between a sceptical and a credulous approach. Moreover, the complex nature of constrained output makes it difficult to handle. This becomes apparent if we consider norm change, like contraction and revision of norms. The constrained input/output logic framework becomes relatively complex and cumbersome. Here, we follow the work of Boella *et al.* [16] and call the generators of unconstrained output *rules*.

3.1 Input/Output Logic

In this section, we give a general introduction to input/output logic. For a deeper look into the input/output logic framework, the reader is referred to the work of Makinson and van der Torre [45] and Parent and van der Torre [57].

A *rule* is a pair of propositional formulas,⁶ called the antecedent and consequent of the rule.

Definition 16 (Rules [42]). *Let L be a propositional logic built on a finite set of propositional atoms A . A rule-based system $R \subseteq L \times L$ is a set of pairs of L , written as $R = \{(a_1, x_1), (a_2, x_2), \dots, (a_n, x_n)\}$.*

Rules allow the derivation of formulas, like the derivation of obligations and prohibitions in a legal code. Which obligations and prohibitions can be derived depends on the factual situation (i.e. the *context* or *input*), which is a propositional formula.

Definition 17 (Operational semantics [42]). *An input/output operation $out : \mathcal{P}(L \times L) \times L \rightarrow \mathcal{P}(L)$ is a function from the set of rule-based systems and contexts to a set of sentences of L .*

Note that operator *out* satisfies the principle of irrelevance of syntax. The simplest input/output logic defined by Makinson and van der Torre is the so-called simple-minded output.

Definition 18 (Simple-minded output [42]). *Proposition x is in the simple-minded output of the set of rules R in context a , written as $x \in out_1(R, a)$, if there is a set of rules $(a_1, x_1), \dots, (a_n, x_n) \in R$ such that $a_i \in Cn(a)$ and $x \in Cn(x_1 \wedge \dots \wedge x_n)$, where $Cn(a)$ is the consequence set of a in L .*

A set of rules is said to ‘imply’ another rule (a, x) if and only if x is in the output in context a .

Definition 19. *Rule ‘implication’ by Makinson and van der Torre [42]] Rule (a, x) is ‘implied’ by rule-based system R , written as $(a, x) \in out(R)$, if and only if $x \in out(R, a)$.*

As Makinson and van der Torre observe, the relation between the ‘implication’ among rules $(a, x) \in out(R)$ and the ‘operational semantics’ $x \in out(R, a)$ has an analogy in classical logic, where the pair $a \models x$ is equivalent to the membership of x in the consequence set of a , written as $x \in Cn(a)$.

Definition 20. [16] *Function out is a closure operation when the following three conditions hold:*

⁶One may also use a first-order, temporal or action logic. The choice of classical propositional logic is intended to stay closer to the AGM theory.

reflexivity: $x \in \text{out}(R \cup \{(a, x)\}, a)$ (in other words, $R \subseteq \text{out}(R)$), and if the context is precisely the antecedent of one of the rules, then the output contains the consequent of that rule.

monotony: $x \in \text{out}(R_1, a)$ implies $x \in \text{out}(R_1 \cup R_2, a)$ (in other words, $\text{out}(R_1) \subseteq \text{out}(R_1 \cup R_2)$), and if the set of rules increases, then no conclusions are lost.

idempotence: if $x \in \text{out}(R, a)$, then for all b , we have $\text{out}(R, b) = \text{out}(R \cup \{(a, x)\}, b)$ (in other words, $\text{out}(R) = \text{out}(\text{out}(R))$), and if x is obligatory in context a , then (a, x) can be added to the rule-based system without changing the output.

Makinson and van der Torre show that their seven input/output logics satisfy the Tarskian properties, and their notion of ‘implication’ among rules is therefore a Tarskian consequence relation, a crucial characteristic to incorporating the AGM construction into the framework of input/output logics.

Definition 21. [42] Let $R(a) = \{x \mid (a, x) \in R\}$, and let v be a classical valuation (maxiconsistent set of propositions) or L . Simple-minded, basic, reusable and basic reusable output are defined as follows:

simple minded: $\text{out}_1(R, a) = \text{Cn}(R(\text{Cn}(a)))$

basic: $\text{out}_2(R, a) = \cap \{\text{out}_1(R, v) \mid a \in v\}$

reusable: $\text{out}_3(R, a) = \cap \{\text{out}_1(R, b) \mid a \in \text{Cn}(b), \text{out}_1(R, b) \subseteq \text{Cn}(b)\}$

basic reusable: $\text{out}_4(R, a) = \cap \{\text{out}_1(R, v) : a \in v \text{ and } \text{out}_1(R, v) \subseteq v\}$

Basic output handles reasoning by cases, and reusable output handles iterated detachment [42]. Moreover, for each input/output logic, a corresponding throughput operator is defined by:

$$\text{out}_i^+(R, a) = \text{out}_i(R \cup \{(b, b) \mid b \in L\}, a).$$

As many of the examples discussed in section 1 have shown, normative change has to handle and solve inconsistencies and incoherencies (on the concept of incoherence, see section 3.2 below) between obligations and permissions as two distinctive kinds of regulative rules.

The implication (or derivation) of obligations from a set O of obligatory regulative rules is given by definition 19. With respect to permissions, it is important

beforehand to distinguish, following Alchourrón [2], between *weak* (or negative) permissions and *strong* (or positive) permissions. In its weak sense, a permission to x in context b is just the absence of a prohibition to x in context b . That is, if we consider a set of obligation rules O , then a permission $\langle a, x \rangle$ is implied by O if and only if $\neg x \notin \text{out}(O, b)$ [44].

In its strong or positive sense, a permission is derived from explicit enactments of obligations as well as permissive rules. The output of a set of explicit permissions is defined below:

Definition 22. [44] *Let O be a set of obligations and let $P \subseteq (L \times L)$ be a set of explicit permissions. Then, $(a, x) \in \text{perm}_i(P, N)$ iff $(a, x) \in \text{out}_i(O \cup Q)$ for some singleton or empty $Q \subseteq P$.*

As we have emphasised in section 1.1 when referring to the problem called *network effects*, some difficulties concerning normative change are related to the combination of constitutive and regulative rules in the normative system.

We may model this problem using input/output logics by making the output of a normative set (possibly joined with the input set) the input of the output operation on the other normative set. It is also possible to combine sets for deriving obligations and explicit permissions.

A typical combination of normative sets is given by the definition or qualification, by a constitutive rule, of a concept present in a regulative rule. For instance, a data protection legislation contains a regulative rule establishing that consent by the data subject (*consent*) is a condition for lawful processing of his/her personal data (*process*). Suppose that a platform processes the personal data of its users without explicit consent, considering that authorisation is implicit unless they explicitly object to that processing (*opt-out* model). If an user of that internet platform has not opted out, would the processing of her personal data be lawful? The answer may be found in a constitutive rule stipulating that only the data user's explicit and written authorisation for processing counts as consent (*opt-in* model). If the set of constitutive rules contain such a rule, then an opt-out model does not count as valid consent for personal data processing. This example of legal reasoning may be modelled by a combination of a set C of constitutive rules and sets O and P of regulative rules, where an output operator on the set of constitutive rules delivers the inputs for the output operator on the sets of regulative rules.

We shall use a general definition of the relation between a constitutive and a regulative rule in a derivation:

Definition 23. *Let $A \subseteq L$, $I \in \{A, \emptyset\}$, let out_i and out_j be output operators, and let C and R be constitutive and regulative sets of rules respectively. Then, the combined output of C and R is defined as:*

$$out_{ij}(C, R, A) = out_i(R, out_j(C, A) \cup I).$$

The definition and the results regarding the contraction operator in section 3.7 covers, with straightforward adaptations, both cases of combinations, i.e. constitutive with permissive rules and constitutive with obligation rules, as follows:

$$\begin{aligned} out_{i,j}(O, C, A) &= out_i(O, out_j(C, A) \cup I); \\ perm_{i,j}(P, C, A) &= perm_i(P, out_j(C, A) \cup I). \end{aligned}$$

In the examples used throughout this article, we shall consider combined *out* and *perm* operators in which $I = A$. To formalise the above example on consent for lawful data processing using a combination of sets of normative rules, let us consider the following normative sets:

$$\begin{aligned} C &= \{(opt-in, consent), (opt-out, \neg consent)\} \\ P &= \{(consent, process)\} \\ O &= \{(\neg consent, \neg process), \} \end{aligned}$$

The normative system implies that $(opt-in, process) \in perm_{1,1}(P, C)$ and that $(opt-out, \neg process) \in perm_{1,1}(O, C)$. That is, it is permitted to process personal data if authorisation was obtained by an opt-in model, while it is forbidden to process that data if the model used was opt-out.

3.2 Consistency and Coherence of Normative Systems

As example 4 in section 1.2 shows, constitutive rules may be responsible for genuine normative conflicts when combined with a regulative set. In order to model this feature, it should be possible to verify regulative sets that are consistent but whose combination with a constitutive set implies inconsistent conditional norms. To avoid confusion, let us qualify regulative sets as consistent or inconsistent and combinations of constitutive sets with regulative sets as coherent or incoherent.

Consistency is defined with respect to a given context. We say that a normative set N is *b-consistent* if and only if $(b, \perp) \notin out(N)$. Accordingly, a combination (C, R) is *b-coherent* if and only if $(b, \perp) \notin out(C, R)$. If we have a set of obligations O and a set of explicit permissions P , then such normative sets are *b-consistent* if and only if for any sentence x , it is not the case that $(b, x) \in perm(O, P)$ and $(b, \neg x) \in out(O)$. Accordingly, a combination of a set of constitutive rules and a set of obligations and the same set of constitutive rules and a set of permissions

is *b-coherent* if and only if, for any sentence x , it is not the case that $(b, x) \in \text{perm}(O, P, C)$ and $(b, \neg x) \in \text{out}(O, C)$.

When should we then consider a normative system to be generally consistent or coherent? We may consider two extreme possibilities for such definitions.

The first extreme would be to consider a normative system *consistent* if it is consistent for all possible inputs, that is, to demand $\perp$ -consistency. This conception would limit the possibility of giving opposite commands in logically independent conditions, since $N = \{(a, x), (b, \neg x)\}$ would be inconsistent.

The other extreme would be to consider a normative system *consistent* if it is consistent for a tautological input, i.e. to demand $\top$ -consistency. This conception also seems inadequate because normative sets with genuine conflicts such as $N = \{(a, x), (a, \neg x)\}$ would be rendered consistent.

As a middle ground, we shall consider a normative set N consistent if it is *b-consistent* for every b such that $b \in Cl(a)$ and $a \in \text{body}(N)$ where $\text{body}(N) = \{b : (b, x) \in N\}$. That is, a normative set is consistent if there is no condition explicitly mentioned in its conditional rules that would, as input, deliver inconsistent outputs. Accordingly, a combination (C, R) is coherent if it is *b-coherent* for every b such that $b \in Cl(a)$ and $a \in \text{body}(C)$.

Therefore, we may have a consistent set R but an incoherent combination (C, R) , which would demand a contraction to restore coherence.

Let us formalise example 4 in the model proposed here. Following Maranhão and de Souza [52], we shall employ a basic reusable output operator (out_4) for the set of constitutive rules, and a basic output operator (out_2) for the sets of regulative (obligatory and permissive) rules. Recall that the example referred to a normative system where the police have the power to access (*acc*) property items (*prop*) in a search & seizure order (*sord*) but are forbidden from accessing ongoing communication (*com*) without an interception order (*iord*). The pertinent question is whether an exchange of messages stored on a mobile phone (*sms*) counts as data (*dat*) or as communication (or both). This normative system could be represented by the following normative sets of constitutive (C), regulative obligation (O) and regulative permission (P) rules:

$$\begin{aligned} C &= \{(sms, com), (sms, dat), (dat, prop)\} \\ P &= \{(prop \wedge sord, acc), (com \wedge iord, acc)\} \\ O &= \{(com \wedge \neg iord, \neg acc), (prop \wedge \neg sord, \neg acc)\} \end{aligned}$$

The corresponding normative theory is both consistent and coherent as there is no explicit condition in these normative sets that can, by itself, deliver a contradiction as output. However, given that a message exchange on a mobile phone

collected during an authorised search is both stored data and a form of ongoing communication, a search & seizure order to check message exchanges would deliver a contradiction, that is, we have both $(sms \wedge sord \wedge \neg iord, acc) \in perm_{2,4}(O, P, C)$ and $(sms \wedge sord \wedge \neg iord, \neg acc) \in out_{2,4}(O, C)$. Hence, the normative system is $(sms \wedge sord \wedge \neg iord)$ -incoherent, and a contraction should take place to restore coherence for that specific context.

There are different ways to reach this goal. And the task of legal interpretation, doctrinal or judicial, is to choose and justify such choices. It is possible to restore coherence by handling the definitions involved, that is, by contracting the set of constitutive rules, by contracting the set of regulative rules, or by deleting rules from both sets. We shall explore these alternatives in section 3.7 below.

3.3 Contraction of Normative Systems

Boella *et al.* [16] defined a rule set as a set of rules closed under an input/output logic ($out(R)$), and generalised the AGM postulates as postulates for the revision of norms. In order to keep an abstract approach and obtain general results without specifying a particular logic, they used operator out to refer to any input/output logic. Operation $out(R) \oplus (a, x)$ indicates the expansion of a rule based-system R by a new rule, operation $out(R) \ominus (a, x)$ denotes the contraction of a rule (a, x) from $out(R)$, and operation $out(R) \circledast (a, x)$ indicates the revision of $out(R)$ by new rule (a, x) .

Like AGM expansion, the definition of rule expansion is straightforward. The new rule that is enforced does not cause any conflict with the existing legal code. Hence, rule (a, x) is added to $out(R)$ together with all the rules that can be derived from the union of $deriv(R)$ and (a, x) : $out(R) \oplus (a, x) = out(R \cup \{(a, x)\})$.

Definition 24. [16] *Let out be an input/output logic. A rule contraction operator $\ominus$ satisfies the following postulates:*

R-1: $out(R) \ominus (a, x)$ is closed under out (closure or type)

R-2: $out(R) \ominus (a, x) \subseteq out(R)$ (inclusion or contraction)

R-3: If $(a, x) \notin out(R)$, then $out(R) = out(R) \ominus (a, x)$ (vacuity or min. action)

R-4: If $(a, x) \notin out(\emptyset)$, then $(a, x) \notin out(R) \ominus (a, x)$ (success)

R-5: If $(a, x) \in out(R)$, then $out(R) \subseteq (out(R) \ominus (a, x)) \oplus (a, x)$ (recovery)

R-6: If $out(\{(a, x)\}) = out(\{(b, y)\})$, then $out(R) \ominus (a, x) = out(R) \ominus (b, y)$ (*extensionality*)

As we have seen in definition 6, the last two AGM postulates ((- 7)-(- 8)) are optional and refer to conjunctions. Since conjunctions are not defined for rules, we restrict ourselves to the basic postulates.

A few words are due about the success postulate. The *success* postulate for rule contraction says that if $x \notin out(\emptyset, a)$, then $x \notin out(R \ominus (a, x), a)$. There are several ways in which a set of rules can be contracted. The purpose of the postulates is to distinguish admissible solutions from inadmissible ones. However, unlike in AGM theory revision, the question here concerns not only what and how much to contract, but also *which inputs* to contract. Boella *et al.* [16] show with the aid of an example that sometimes, in order to obtain a rule-based system that satisfies the success postulate, one needs to *add* some rules.

Another issue is the characterisation of the minimal rule contraction operators. We have seen that in AGM, one interpretation of the postulates is to impose the economical principle. That is, when performing a rule contraction operator, we want to keep as much as possible. However, a syntactic characterisation of minimal rule contraction encounters some problems. In AGM, thanks to the closure postulate (i.e. belief sets are closed under consequence), if $y \notin (K - x)$, then we also have that $x \wedge y \notin (K - x)$. Likewise, if $(a, x) \notin out(R) \ominus (a, x)$, then also $(a, x \wedge y) \notin out(R) \ominus (a, x)$. However, this is not the only consequence of the success postulate for rule contraction. For example, for all six input/output logics considered here, if $(a, x) \notin out(R) \ominus (a, x)$, then also $(a \vee b, x) \notin out(R) \ominus (a, x)$.

Other logical relations depend on the input/output logic used. For example, for basic output out_2 , if $(a, x) \notin out(R) \ominus (a, x)$, then we have either $(a \wedge b, x) \notin out(R) \ominus (a, x)$ or $(a \wedge \neg b, x) \notin out(R) \ominus (a, x)$. In other words, if $(a, x) \notin out(R) \ominus (a, x)$ and $(a \wedge b, x) \in out(R) \ominus (a, x)$, then $(a \wedge \neg b, x) \notin out(R) \ominus (a, x)$. These relations do not hold for simple-minded output out_1 . Likewise, a similar property based on the inverse of CTA holds for reusable output out_3 .

The recovery postulate states that contracting a rule-based system by (a, x) and then expanding by the same (a, x) should leave $out(R)$ unchanged. We will see that such a postulate turns out to be problematic for rule contraction.

Boella *et al.* [16] show that the five postulates considered so far are consistent only for some input/output logics, and not for others. In particular, if we adopt output out_1 or out_3 , then there is no single

Proposition 25. [16]

(R-1) to (R-5) cannot hold together for out_1 or out_3 , but they can hold together for out_2 .

We now turn to the postulates for rule revision.

3.4 Revision of Normative Systems

As in rule contraction, we consider only the first six AGM revision postulates and the rule revision postulates.

Definition 26. [16] Let out be an input/output logic, and $deriv(R)$ a set of rules closed under out . A rule revision operator $\circledast$ satisfies the following postulates:

R $\circledast$ 1: $out(R) \circledast (a, x)$ is closed under out (closure or type)

R $\circledast$ 2: $(a, x) \in (out(R) \circledast (a, x))$ (success)

R $\circledast$ 3: $out(R) \circledast (a, x) \subseteq out(R) \oplus (a, x)$ (inclusion)

R $\circledast$ 4: If $(a, \neg x) \notin out(R \cup (a, x))$ then $out(R) \oplus (a, x) = out(R) \circledast (a, x)$ (vacuity)

R $\circledast$ 5: $(a, \neg x) \in out(R) \circledast (a, x)$ iff $(a, \neg x) \in out(\emptyset)$ (triviality)

R $\circledast$ 6: If $out(\{(a, x)\}) = out(\{(b, y)\})$, then $out(R) \circledast (a, x) = out(R) \circledast (b, y)$ (extensionality)

As seen in section 2, the Levi Identity defines revision $K * A$ as a sequence of a contraction and an expansion. We have seen the correctness of such a definition in observations 10 and 12.

It is worth noting that the controversial recovery postulate (– 5) was not used in observation 12. Boella *et al.* [16] show that the same result can be proven for rule change.

Theorem 27. [16] Given a rule contraction operator, we can define a rule revision operator via the Levi Identity:

$$out(R) \circledast (a, x) = (out(R) \ominus (a, \neg x)) \oplus (a, x).$$

When operator $\ominus$ satisfies rules (R-1) to (R-4) and (R-6), then operator $\circledast$ satisfies rules (R*1) – (R*6).

Not only can belief revision be defined in terms of belief contraction operators, belief contractions can also be defined in terms of belief revisions using the Harper and Levi Identities introduced in section 2 .

However, as recalled in proposition 25, for out_1 and out_3 the revision postulates are consistent and the contraction postulates are not. Thus, a result like observation 10 for normative change does not hold.

We recall from section 2 that the Levi and Harper Identities have been shown to be interchangeable in AGM theory. So, even though there is no theorem corresponding to observation 10 in the general case, one may want to check whether $out(R) \circledast (a, x) = (out(R) \cap out(R) \circledast (a, x)) \oplus (a, x)$ is a consequence of the basic postulates for rule revisions, and whether $out(R) \ominus (a, x) = out(R) \cap ((out(R) \ominus (a, x)) \oplus (a, \neg x))$ can be proven from the basic set of postulates for rule contractions (including the recovery postulate). Boella *et al.* [16] show that the answer to the first question is positive:

Proposition 28. [16] $out(R) \circledast (a, x) = (out(R) \cap out(R) \circledast (a, x)) \oplus (a, x)$.

However, $out(R) \ominus (a, x) = out(R) \cap ((out(R) \ominus (a, x)) \oplus (a, \neg x))$ does not hold in general, i.e. it cannot hold for output out_1 or out_3 .

3.5 Contraction of Normative Bases

Models of belief contraction and revision are built in order to satisfy the demand for minimal change to keep a theory consistent. As we have seen in section 2.1 above, there are two basic strategies for reaching this goal with the syntactic approach. The first consists in selecting the resulting contraction or revision among maximal consistent subsets of the original. The second consists in making an “incision” in the minimal subsets of the theory or base that derived the sentence to be deleted or revised. We shall now follow the second strategy, calling those minimal subsets “arguments”, which are here the base of normative entailments from the set of rules. The construction proceeds basically by making minimal withdrawals from those arguments:

Definition 29. (Argument) $X \subseteq L \times L$ is an argument for (a, x) based on a normative set N if and only if:

- (i) $X \subseteq N$;
- (ii) $(a, x) \in out(X)$;
- (iii) if $X' \subset X$, then $(a, x) \notin out(X')$.

$Args_N(a, x)$ is the set of arguments for (a, x) based on N .

Definition 30. An incision σ is a choice-like function on $\text{Args}_N(a, x)$ to $\wp(L \times L)$ such that:

- (i) $\sigma(\text{Args}_N(a, x)) \subseteq \bigcup \text{Args}_N(a, x)$;
- (ii) $\sigma(\text{Args}_N(a, x)) \cap X \neq \emptyset$, for all $X \in \text{Args}_N(a, x)$.

Definition 31. Let N be a normative set and (a, x) a conditional norm. Then, the contraction of N by (a, x) is defined as:

$$N -_{\sigma} (a, x) = N \setminus \sigma(\text{Args}_N(a, x)).$$

The contraction of a normative set N by a conditional rule (a, x) may also be defined by postulates on a contraction function, as follows.

Definition 32. The contraction of a normative set N by a conditional rule (a, x) is a function $N - : L \times L \rightarrow \wp(L \times L)$ satisfying the following postulates:

- N-1:** if $(a, x) \notin \text{out}(\emptyset)$, then $(a, x) \notin \text{out}(N - (a, x))$ (success)
- N-2:** $N - (a, x) \subseteq N$ (inclusion)
- N-3:** if $(b, y) \in N \setminus N - (a, x)$, then there is $N' \subset N$ such that $(a, x) \notin \text{out}(N')$, but $(a, x) \in \text{out}(N' \cup \{(b, y)\})$ (core-retainment)
- N-4:** if for all $N' \subseteq N$, $(a, x) \in \text{out}(N')$, if and only if $(b, y) \in \text{out}(N')$, then $N - (a, x) = N - (b, y)$ (uniformity)

The representation theorem below is easily adapted from Hansson's representation theorem for base contraction (observation 15):

Theorem 33. $N -_{\sigma} (a, x) = N - (a, x)$.

3.6 Refinement of Normative Bases

As we have noticed above for output operators stronger than basic output out_2 , the following property holds: if $(a, x) \notin \text{out}(R) \ominus (a, x)$, then either $(a \wedge b, x) \in \text{out}(R) \ominus (a, x)$ or $(a \wedge \neg b, x) \in \text{out}(R) \ominus (a, x)$. Hence, in every contraction of a conditional obligation (a, x) from a closed normative set R , based on an underlying logic at least as strong as basic output, the resulting contracted set $\text{out}(R) \ominus (a, x)$ will include a “weakened” version of the conditional, that is, either $(a \wedge b, x)$ or $(a \wedge \neg b, x)$. It is possible to specify in the selection function which weakened version shall remain. This was the basic intuition underlying the operator called *refinement* proposed by

Maranhão [47], which was aimed at modelling the introduction of exceptions to rules by legal interpretation. For instance, given a normative system that delivers absolute prohibition of abortion, $(\top, \neg \text{abort}) \in O$, a defence of abortion in the case of an anencephalic foetus would not be a proposal for permitting abortion in any context. Hence, the contraction of $(\top, \neg \text{abort})$ from that system should make reference to that specific exception, which means that in the absence of anencephaly, abortion should remain forbidden in that normative system, in the name of minimal change. That is, $(\neg \text{anenceph}, \neg \text{abort})$ should still be derivable from normative system O , while the prohibition should cease to hold in the exceptional case, that is $(\text{anenceph}, \neg \text{abort}) \notin \text{out}(O)$.

By specifying the exception in the selection function, this result follows from the principle of minimality if the normative set is closed and the logic is at least as strong as a basic output. However, for normative bases (not closed sets), deleting $(\text{anenceph}, \neg \text{abort})$ from the set of consequences of O would be tantamount to excluding $(\top, \neg \text{abort})$ from normative set O , and therefore $(\neg \text{anenceph}, \neg \text{abort})$ would not be derived anymore.

But it is possible to define a refinement as a particular case of a *conservative contraction* [49]. That is, it expands the normative set with rules that are entailed by the rule to be contracted, and which include the exceptional factor and its negation in the antecedent.

Definition 34. (*Refinement*) Let $f \in L$, N be a normative system and let $(a, x) \in \text{out}(N)$, where out is at least as strong as a basic output. Then, the refinement of N and (a, x) by factor f is $N \otimes^f (a, x) = N^* -_{\theta_{N^*}} (a, x)$ where $N^* = N \cup \{(f \wedge a, x), (\neg f \wedge a, x)\}$ and $(\neg f \wedge b, y) \notin \theta(\text{Args}_{N^*}(a, x))$. We call factor f an *exception* to (a, x) in the resulting refined normative system.

Proposition 35. *The refinement operator satisfies the following success properties:*

- $(a, x) \notin N \otimes^f (a, x);$
- $(a, x), (f \wedge a, x) \notin N \otimes^f (a, x);$
- $(\neg f \wedge a, x) \in N \otimes^f (a, x).$

3.7 Contraction of Combined Normative Bases

As we have seen in section 3.2, the combination of a constitutive set of rules and regulative sets of permissions and obligations may give rise to genuine incoherencies, that is, the delivery of incompatible rulings, even though the sets of obligations and permissions are consistent. This happens when a given input activates definitions in

the constitutive set that triggers logically independent rules with conflicting outputs. As we have suggested, restoring coherence would involve deciding between several alternatives that may change the set of constitutive rules, or the set of regulative rules, or both. In this section, we are going to introduce a formal framework for the operation of contracting normative systems that combine sets of constitutive rules (which we shall call a *constitutive set*) and regulative rules (which we shall call a *regulative set*).

For $A \subseteq L$, output operators out_i and out_j , constitutive set C and regulative set R , we shall use the following conventions:

- (i) $out_i(C, R, A)$ if $i = j$;
- (ii) $out_{ij}(C, R, a)$ denoting $out_{ij}(C, R, \{a\})$;
- (iii) $(a, x) \in out_{ij}(C, R)$ if $x \in out_{ij}(C, R, a)$.

We call the pair of normative sets (C, R) the combination of C and R or the combination (C, R) .

Below, we build and characterise operators to perform the three kinds of changes in normative systems that combine constitutive and regulative rules. The first operator, called *constitutive contraction*, contracts only the constitutive set. The second operator, called *regulative contraction*, contracts the regulative set. The *combined contraction* operator may contract both in order to delete a norm from the combination of the constitutive and regulative sets.

Definition 36. (*Constitutive contraction*) The constitutive contraction of a combination (C, R) by a conditional norm (a, x) is a function $C -_R : L \times L \rightarrow \wp(L \times L)$ satisfying the following postulates:

- C-1:** if $(a, x) \notin out_i(\emptyset, R)$, then $(a, x) \notin out_i(C -_R(a, x), R)$ (success)
- C-2:** $C -_R(a, x) \subseteq C$ (inclusion)
- C-3:** if $(b, y) \in C \setminus C -_R(a, x)$, then there is $C' \subset C$ such that $(a, x) \notin out_i(C', R)$, but $(a, x) \in out_i(C' \cup \{(b, y)\}, R)$ (core-retainment)
- C-4:** if for all $C' \subseteq C$ it is the case that $(a, x) \in out_i(C', R)$ if and only if $(b, y) \in out_i(C', R)$, then $C -_R(b, y) = C -_R(a, x)$ (uniformity)

Definition 37. (*Regulative contraction*) The regulative contraction of a combination C, R by a conditional norm (a, x) is a function $R -_C : L \times L \rightarrow \wp(L \times L)$ satisfying the following postulates:

- R-1:** if $(a, x) \notin \text{out}_i(C, \emptyset)$, then $(a, x) \notin \text{out}_i(C, R -_C(a, x))$ (success)
- R-2:** $R -_C(a, x) \subseteq R$ (inclusion)
- R-3:** if $(b, y) \in R \setminus R -_C(a, x)$, then there is an $R' \subset R$ such that $(a, x) \notin \text{out}_i(C, R')$, but $(a, x) \in \text{out}_i(C, R' \cup \{(b, y)\})$ (core-retainment)
- R-4:** if for all $R' \subseteq R$, $(a, x) \in (C, R')$ if and only if $(b, y) \in \text{out}_i(C, R')$, then $R -_C(a, x) = R -_C(b, y)$ (uniformity)

We use the following conventions for the definition of the combined contraction of normative sets:

- (i) if $(C, R) - (a, x) = (C^-, R^-)$, then $(C, R) \setminus (C, R) - (a, x) = (C \setminus C^-, R \setminus R^-)$;
 (ii) $\bigcup(C, R) = \bigcup\{C, R\}$.

Definition 38. (Combined contraction) The combined contraction of the combination (C, R) by a conditional norm (a, x) is a function $(C, R) - : L \times L \rightarrow \wp(L \times L) \times \wp(L \times L)$ satisfying the following postulates:

- C/R-1:** if $(a, x) \notin \text{out}_i(\emptyset)$, then $(a, x) \notin \text{out}_i((C, R) - (a, x))$ (success)
- C/R-2:** if $(C, R) - (a, x) = (C^-, R^-)$, then $C^- \subseteq C$ and $R^- \subseteq R$ (inclusion)
- C/R-3:** if $(b, y) \in \bigcup(C, R) \setminus (C, R) - (a, x)$, then there is a $C' \subseteq C$ and $R' \subseteq R$ such that $(a, x) \notin \text{out}_i(C', R')$, but $(a, x) \in \text{out}_i(C' \cup \{(b, y)\}, R')$ or $(a, x) \in \text{out}_i(C', R' \cup \{(b, y)\})$ (core-retainment)
- C/R-4:** if for all $C' \subseteq C$ and $R' \subseteq R$, it is the case that $(a, x) \in \text{out}_i(C', R')$ if and only if $(b, y) \in \text{out}_i(C', R')$, then $(C, R) - (a, x) = (C, R) - (b, y)$ (uniformity)

Now we will define a general construction for kernel contraction of combined normative sets, from which we may specify constitutive, regulative and combined contraction operators.

Definition 39. (Combined argument) A combination (X, Y) is a combined argument for (a, x) based on the combination (C, R) of a constitutive set C and a regulative set R if and only if:

- (i) $X \subseteq C$;
 (ii) $Y \subseteq R$;
 (iii) $(a, x) \in \text{out}_i(X, Y)$;
 (iv) if $X' \subset X$, then $(a, x) \notin \text{out}_i(X', Y)$;
 (v) if $Y' \subset Y$, then $(a, x) \notin \text{out}_i(X, Y')$.

We denote by $Args_{(C,R)}(a, x)$ the set of combined arguments for (a, x) based on (C, R) . Now we will define the incision function for choosing rules from the minimal arguments delivering the rule to be excluded.

Definition 40. *An incision is a choice-like function on $Args_{(C,R)}(a, x)$ to $\wp(L \times L)$ such that:*

- (i) if $Args_{(C,R)}(a, x) = \{(X_i, Y_i) : i \in I\}$,
then $\sigma(Args_{(C,R)}(a, x)) \subseteq \bigcup_{i \in I} (X_i \cup Y_i)$;
- (ii) $\sigma(Args_{(C,R)}(a, x)) \cap (X_i \cup Y_i) \neq \emptyset$ for every $(X_i, Y_i) \in Args_{(C,R)}(a, x)$.

The general definition encompasses incisions that choose rules from both normative sets at the same time, incisions that choose only regulative rules, and incisions that choose only constitutive rules. The definitions above restrict the incision functions to choosing only constitutive rules or only regulative rules.

Definition 41. *An incision on $Args_{(C,R)}(a, x)$ is constitutive if and only if $\sigma(Args_{(C,R)}(a, x)) \cap R = \emptyset$.*

Definition 42. *An incision on $Args_{(C,R)}(a, x)$ is regulative if and only if $\sigma(Args_{(C,R)}(a, x)) \cap C = \emptyset$.*

Now we will use a general definition for contraction based on the incision function. Of course, if we use a constitutive incision, the result will be a constitutive contraction. Similarly, if we use a regulative incision, the result will be a regulative contraction.

Definition 43. *(Contraction) Let (C, R) be a combination of normative sets and (a, x) a conditional norm. Then, the contraction of (C, R) by (a, x) based on incision σ is defined as $(C, R) -_{\sigma} (a, x) = (C^-, R^-)$ where $C^- = C \setminus \sigma(Args_{(C,R)}(a, x))$ and $R^- = R \setminus \sigma(Args_{(C,R)}(a, x))$.*

The theorems below show that the postulates for constitutive, regulative and general contraction characterise the respective constructions.

Theorem 44. *[52] A contraction of (C, R) by (a, x) based on a constitutive incision σ is a constitutive contraction, that is, $(C, R) -_{\sigma} (a, x) = (C -_R(a, x), R)$. Moreover, given a constitutive contraction, there is a constitutive incision σ such that $(C, R) -_{\sigma} (a, x) = (C -_R(a, x), R)$.*

Theorem 45. *[52] A contraction of (C, R) by (a, x) based on a regulative incision σ is a regulative contraction, that is, $(C, R) -_{\sigma} (a, x) = (C, R -_C(a, x))$. Moreover, given a regulative contraction, there is a regulative incision σ such that $(C, R) -_{\sigma} (a, x) = (C, R -_C(a, x))$.*

Theorem 46. [52]

$$(C, R) -_{\sigma} (a, x) = (C, R) - (a, x).$$

The contraction operators discussed here do not involve constraints on the choice of incision function that will determine the result of the contraction operation. Therefore, there is no preference for a regulative contraction over a constitutive or combined contraction.

This feature may be illustrated by example 4, which was formalised in section 3.2. In that case, a contraction to avoid $sms \wedge sord \wedge \neg iord$ -incoherent would have the following alternatives for the incisions: $(C, O) -_{\sigma} (sms \wedge sord \wedge \neg iord, \neg acc)$ or $(C, P) -_{\sigma} (sms \wedge sord \wedge \neg iord, acc)$, each of which is determined by any of the following unitary incision functions: $\sigma_1 = \{(sms, dat)\}$, or $\sigma_2 = \{(sms, com)\}$, or $\sigma_3 = \{(data, prop)\}$, or $\sigma_4 = \{(prop \wedge sord, acc)\}$, or $\sigma_5 = \{(com \wedge \neg iord, \neg acc)\}$.

The controversy within the Brazilian Superior Court of Justice discussed in section 1.2 involved two of these alternative contractions. The first decision was a constitutive contraction based on σ_1 , where the court contended that message exchanges are communications in flux, which demanded a specific order to intercept the conversation.

In turn, the second decision by the Brazilian court was a conservative contraction based on σ_2 , contending that message exchanges should not be considered as ongoing communication. The same alternative contraction was chosen by the German court. The underlying reason for these choices was the weight given to the constitutional value of freedom of communication, which is demoted by such access to the content of an individual's mobile phone. The demotion of freedom of communication was considered stronger than the demotion of property rights. Hence, the association of "text messaging" with "stored data" and, therefore, with "property" (instead of its association with "personal communication") coheres with an underlying valuation where property rights are outweighed by public safety concerns. The German decision also involved a concern about the constitutional right of informational autonomy as the core of data protection. According to the court's argumentation, this right was not violated because the data subject could have destroyed the data in her possession.

Notice that both courts decided not to revise the regulative rules, only stipulate the conceptual qualification of text messaging. The contraction of the regulative set would be inadequate. The first alternative contraction, σ_4 , would lead to the absence of an explicit authorisation to search property items, while the other alternative contraction, σ_5 would exclude the prohibition to intercept communications. Nevertheless, the court could have considered less intrusive interventions on the set

of regulative rules by, for instance, treating the case of text messaging as an exception to search orders on data. That is, in order to reach a coherent normative system in that context (to avoid $sms \wedge sord \wedge \neg iord$ -incoherence), the court could have refined the set of obligations, which in the model would be represented by a refinement operator ensuring that $(\neg sms \wedge prop \wedge sorder, acc) \in P \otimes^{sms} (prop \wedge sorder, acc)$.

The resulting contraction would then be either constitutive or regulative. However, there can be genuine combined contractions on sets of constitutive and regulative rules. Consider, for instance, a variation on example 4, where an order to investigate an individual (*order*) would encompass both a search & seizure procedure and the interception of any communication. We would have the following sets in the normative system:

$$C = \{(sms, dat), (data, prop), (sms, com)\}$$

$$P = \{(com \wedge order, acc), (prop \wedge order, acc)\}$$

According to that normative system, police officers are authorised to access the content of the message exchange stored on the cell phone with a general order authorising the investigation of an individual. Now suppose that the legislator derogates from the positive permission to access the content of text messages stored on a mobile phone, or that legal interpretation (judicial or doctrinal) considers such a permission to be unconstitutional for violating the fundamental right to privacy. In that case, a contraction $(C, P) - (sms \wedge order, acc)$ involves choosing from the following incisions:

$$\sigma_1 = \{(sms, dat), (com, acc)\}$$

$$\sigma_2 = \{(dat, prop), (com, acc)\}$$


$$\sigma_3 = \{(sms, com), (prop, acc)\}$$

$$\sigma_4 = \{(sms, com), (sms, dat)\}$$

$$\sigma_5 = \{(sms, com), (dat, prop)\}$$

$$\sigma_6 = \{(prop, acc), (com, acc)\}$$

The contractions based on σ_{1-3} are combined contractions, while those based on σ_{4-5} are constitutive contractions. The contraction based on σ_6 is the only alternative based on regulative contraction. The figure below illustrates incision σ_1 , where each dash linking two nodes is a pair, and each node is proposition:



The contraction based on σ_1 is indeed the most reasonable choice. A regulative contraction is clearly undesirable, since it would make any search unauthorised, resulting in a normative system that completely disregards the value of public safety and containing useless definitions. On the other hand, a constitutive contraction based on σ_4 would not address the crucial question in this case, which is how to legally qualify text messaging. In turn, the constitutive contraction σ_5 would make it impossible to search for any document on the premises, in spite of defining that text messaging counts as data. The combined contraction σ_2 would be similar to σ_1 , with the effect of favouring freedom of communication over public safety. However, it would also have the undesirable effect of hindering access to any data in a search procedure. For a similar reason, σ_3 would be inadequate with regard to the intuition that the protection of property rights has less weight than the protection of freedom of communication when balancing public safety concerns.

4 Challenges and Open Problems with the AGM Approach

In this final section, we will discuss some open problems and relevant questions that are the object of mainstream research on normative change with the AGM approach.

As we have seen in section 2.2, one of the main challenges and criticisms of the AGM approach is the potential indeterminacy of the result of a contraction, revision or refinement of the normative system, which depends on choices about the proper selection or incision functions to determine the result. This feature is sometimes seen as a disadvantage compared to the syntactic approach, where the syntax of a particular rule is the object of change.

Actually, as we have argued in section 1.1, what we have called the “indeterminacy problem” is not really a defect of the representational model, but is a real

feature of legal reasoning about normative change that should be captured by the model itself. As a representation of the activity of legal interpretation, it is particularly interesting to show what are the alternative interpretations for different acts of the derogation, making it clear that a particular interpretation involves choices.

Although there may be some alternative interpretations that are clearly inadequate and would be immediately rejected by a jurist, it is important to investigate the criteria for rejection and represent them in the model. It is also a fact that there may be a doctrinal or judicial controversy concerning the defensible results of a normative change, as illustrated in example 4, and we believe that the model should be able to express these different available choices as an adequate representation of legal reasoning. So we see the indeterminacy reflected in the model as an advantage of the AGM approach.

However, there is also an *onus* on this model to provide criteria that would reflect the consensual choices (in the sense of consensus on action, not consensus on explicit convention) reached by legal practitioners and jurists on normative change. Hence, one of the main challenges to research on normative change is to find and model criteria for determining rational choices from alternative normative systems resulting from change operations.

When discussing the examples formalised in section 3.7, we provided some reasons for preferring certain incisions over others. The arguments used there to justify the choice of a particular incision were all domain-specific. Nevertheless, the discussion provided at least two important clues for developing more abstract constraints.

The first clue is related to Makinson and van der Torre's discussion on constraints for I/O-logics [43] suggesting a distinction between rule maximisation (*maxrule*: maximising the preservation of rules in order to satisfy a constraint) and output maximisation (*maxout*: maximising the preservation of outputs in order to satisfy a constraint). The Mobius Strip example is a radical case and may be seen as a contraction. Consider $N = \{(\top, a), (a, b), (b, \neg a)\}$. The contraction $N - (\top, \perp)$ has two possible outcomes: $N_1 = \{(\top, a)\}$ or $N_2 = \{(a, b), (b, \neg a)\}$. While N_1 satisfies *maxout* and fails *maxrule*, N_2 satisfies *maxrule* and fails *maxout*.

Indeed, constitutive contractions tend to favour *maxrule* and sacrifice *maxout*, since intermediary concepts may be connected to different rules. As we have indicated in section 1.1, the network effects problem regarding normative change alerts us that suppressing relevant connections between normative concepts may render regulative rules inapplicable, while deleting regulative rules may change our understanding of some normative concepts. The construction of the contraction operators for combined normative sets in this article was based on rule maximisation, but future investigations should try to find reasonable constraints to temper the demand for *maxrule* with the demand for *maxout*.

The second clue is the role of values that drive the choices among possible outcomes of a change function. The positively enacted rules (constitutive and regulative) on which the legal order are built are the outcomes of (legislative or judicial) deliberations on relevant societal values (moral considerations, political goals, fundamental rights). Those societal values inform the interpretation of authoritative decisions in the application of the rules of the normative system when evaluating the legality of actions in particular contexts. Such values may be considered as external to the normative system or as internal to it in the form of constitutional rights and principles. Thus, if one conceives of legal interpretation as a dynamic of normative change, as suggested in section 1.2, then enriching the model with reasoning about balancing values would provide relevant criteria for choosing between the resulting contracted, revised or refined normative systems, a line of research recently pursued by Maranhão [50] and Maranhão and Sartor [53].

If one takes seriously the representation of legal interpretation as normative change, and succeeds in modelling relevant criteria for choosing among possible systems resulting from contractions, revisions and refinements, then argumentation frameworks could be developed to model argumentation by legal doctrine to determine the best interpretation. That is, there could be a model of argumentation about the results of normative change. Such an argumentation process would put forward defeasible arguments about competing goals of legal interpretation (consistency, coherence with underlying political morality, completeness, precision, adherence to positively enacted rules and natural language, etc.).

The incorporation of tools to represent reasoning about values in the model of normative change will inevitably lead to the need to adapt the change functions to non-monotonic logic, including input/output logics where its rules are default (see [56]). There is a fairly dominant trend in legal theory [9] and in the literature of artificial intelligence & law (see [12], [61] and [69]) of considering reasoning about values as defeasible, where consideration of additional values in a particular context may defeat reasons for particular actions in a framework of an overall appreciation of those relevant values. Hence, as already mentioned in section 2.2, the AGM methodology should be adapted to systems with underlying logics that are not monotonic, as pursued recently by Zhuang *et al.* [79], Casini and Meyer [20], and Casini *et al.* [19]. Since the addition of new values or considerations related to values may defeat some implications or reasons for action, with the AGM approach such systems will reflect an aspect of the syntactic approach where a “contraction” is obtained by adding rules to the normative set [50]. As argued in sections 1.2 and 1.4, the representation of legal interpretation should involve values, and that aspect may point to incorporating methods of revision provided by what we have called the preferential approach.

Another important observation concerning applications of the formal models of normative change, emphasised in sections 1.2 and 1.4, is how to adequately represent the two dimensions of normative change: the dimension of validity, which we believe is better reflected by the AGM approach, and the dimension of efficacy, which seems to be better captured by the syntactic approach. Integrating both perspectives would also demand formal comparisons between these approaches. Where the AGM approaches focus on changes in the normative system, it is pertinent to ask whether and how the resulting system can be captured by syntactic modifications of the rules and how alternative interpretations can be represented. Where the syntactic approaches focus on the syntactical representation of the time span of the efficacy of rules and how to block or enable their effects, it is pertinent to ask whether the enabled rules in a given time span can be represented by a temporal dynamic for subsystems of the whole system of valid rules (containing the rules that are enabled at a given period). Efforts to enrich the syntactic representation of rules within the AGM approach with, for instance, time labels [74], are also important for modelling reasoning that closely reflects real-life examples of the complex interaction between the period of a rule's efficacy and the time span of its validity in the legal system.

There are also conceptual and formal results to be pursued by researchers working on the AGM approach. For instance, there are still no formal characterisations of revision and refinement for changing combined normative sets. It is also relevant to explore constructions of revision from contraction and vice versa for some input/output logics where the Harper and Levi Identities would not hold (see section 3.4). A general theory of revision functions on different sorts of architectures of input/output logics (combinations of normative sets within the input/output logics framework) would also be a relevant theoretical achievement to ground future research of applications that explore particular architectures [17, 53] for more complex architectures).

The constructions discussed in this article were based on original input/output logics (simple-minded, basic, reusable and basic reusable) introduced by Makinson and van der Torre [42]. It would be interesting to apply the AGM approach to input/output logics with constraints [43] and other variants [58, 59].

Acknowledgments

Leendert van der Torre acknowledges financial support from the Fonds National de la Recherche Luxembourg (INTER/Mobility/19/13995684/DLAI/van der Torre). Giovanni Casini has been supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215.

References

- [1] Aulis Aarnio. *On Legal Reasoning*. Turun Yliopisto, 1977.
- [2] Carlos E Alchourrón. Logic of Norms and Logic of Normative Propositions. *Logique et analyse*, 12:242–268, 1969.
- [3] Carlos E. Alchourrón. Detachment and Defeasibility in Deontic Logic. *Studia Logica*, 57:5–18, 1996.
- [4] Carlos E Alchourrón and Eugenio Bulygin. *Normative Systems*. Springer, 1971.
- [5] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [6] Carlos E. Alchourrón and David Makinson. *Hierarchies of Regulations and their Logic*, pages 125–148. Springer Netherlands, 1981.
- [7] Carlos E. Alchourrón and David Makinson. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.
- [8] Robert Alexy. On the structure of legal principles. *Ratio Juris*, 13:294–304, 2000.
- [9] Robert Alexy. Constitutional Rights, Balancing, and Rationality. *Ratio Juris*, 16:131–140, 2003.
- [10] Amalia Amaya. *The Tapestry of Reason: an inquiry into the nature of coherence and its role in legal argument*. Hart Publishing, 2015.
- [11] Aharon Barak. *Purposive Interpretation in Law*. Princeton University Press, 2005.
- [12] Trevor J. M. Bench-Capon and Giovanni Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150:97–142, 2003.
- [13] Donald H. Berman and Carole D. Hafner. Representing teleological structure in case-based reasoning: The missing link. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law (ICAIL)*, pages 50–9. ACM, 1993.
- [14] Norberto Bobbio. Le bon législateur. *Logique & Analyse*, 14(53-54):243–249, 1971.
- [15] Guido Boella, Guido Governatori, Antonino Rotolo, and Leendert van der Torre. Lex minus dixit quam voluit, lex magis dixit quam voluit: A formal study on legal compliance and interpretation. In *AICOL-I/IVR-XXIV’09 Proceedings of the 2009 international conference on AI approaches to the complexity of legal systems: complex systems, the semantic web, ontologies, argumentation, and dialogue*, pages 162–183. Springer, 2010.
- [16] Guido Boella, Gabriella Pigozzi, and Leendert van der Torre. AGM contraction and revision of rules. *Journal of Logic, Language and Information*, 25(3-4):273–297, 2016.
- [17] Guido Boella and Leendert van der Torre. A Logical Architecture of a Normative System. In *Deontic Logic and Artificial Normative Systems - DEON 2006*, pages 24–35. Springer, 2006.
- [18] Guido Boella and Leendert van der Torre. A logical architecture of a normative system. In *Deontic Logic and Artificial Normative Systems - DEON 2006*, pages 24–35. Springer, 2006.

- [19] Giovanni Casini, Eduardo Fermé, Thomas Meyer, and Ivan Varzinczak. A semantic perspective on belief change in a preferential non-monotonic framework. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018.*, pages 220–229, 2018.
- [20] Giovanni Casini and Thomas A. Meyer. Belief change in a preferential non-monotonic framework. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 929–935, 2017.
- [21] James P. Delgrande and Pavlos Peppas. Belief revision in Horn theories. *Art. Int.*, 218:1 – 22, 2015.
- [22] Julie Dickson. Interpretation and Coherence in Legal Reasoning. In Edward N Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, 2016.
- [23] Ronald M Dworkin. *Law’s Empire*. Kermode, 1986.
- [24] Eduardo L. Fermé. On the logic of theory change: Contraction without recovery. *Journal of Logic, Language and Information*, 7(2):127–137, Apr 1998.
- [25] Eduardo L. Fermé and Sven O. Hansson. *Belief Change - Introduction and Overview*. Springer Briefs in Intelligent Systems. Springer, 2018.
- [26] Lon L Fuller. Positivism and Fidelity to Law - A Reply to Professor Hart. *Harvard Law Review*, 71:630–672, 1958.
- [27] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
- [28] Peter Gärdenfors and Hans Rott. Belief revision. In C.J. Hogger D.M. Gabbay and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming. Volume IV: Epistemic and Temporal Reasoning*, pages 35–132. Oxford University Press, 1995.
- [29] Guido Governatori, Francesco Olivieri, Matteo Cristani, and Simone Scannapieco. Revision of defeasible preferences. *Int. J. Approx. Reason.*, 104:205–230, 2019.
- [30] Guido Governatori and Antonino Rotolo. Changing legal systems: legal abrogations and annulments in defeasible logic. *Logic Journal of IGPL*, 18(1):157–194, 2010.
- [31] Guido Governatori, Antonino Rotolo, Francesco Olivieri, and Simone Scannapieco. Legal contractions: a logical analysis. In Enrico Francesconi and Bart Verheij, editors, *ICAIL*, pages 63–72. ACM, 2013.
- [32] Davide Grossi, John-Jules Ch. Meyer, and Frank Dignum. The many faces of counts-as: A formal analysis of constitutive rules. *Journal of Applied Logic*, 6:192–217, 2008.
- [33] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [34] Sven O. Hansson. Reversing the Levi identity. *Journal of Philosophical Logic*, 22:637–669, 1993.
- [35] Sven O. Hansson. *A Textbook of Belief Dynamics: Theory Change and Database Up-*

- dating*. Kluwer Academic Publishers, 1999.
- [36] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, KR'91, pages 387–394, San Francisco, CA, USA, 1991. Morgan Kaufmann Publishers Inc.
 - [37] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 3(52):263–294, 1991.
 - [38] Hans Kelsen. Derogation. In *Essays in Legal and Moral Philosophy*, pages 261–275. Springer Netherlands, 1973.
 - [39] Sebastien Konieczny and Ramon Pino Perez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 10 2002.
 - [40] Isaac Levi. Subjunctives, dispositions and chances. *Synthese*, 34:423–455, 1977.
 - [41] David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16(4):383–394, Nov 1987.
 - [42] David Makinson and Leendert van der Torre. Input-output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.
 - [43] David Makinson and Leendert van der Torre. Constraints for input-output logics. *Journal of Philosophical Logic*, 30(2):155–185, 2001.
 - [44] David Makinson and Leendert van der Torre. Permissions from an input-output perspective. *Journal of Philosophical Logic*, 32(4):391–416, 2003.
 - [45] David Makinson and Leendert van der Torre. What is input/output logic? In L'owe Benedikt, Malzkorn Wolfgang, and R'asch Thoralf, editors, *Foundations of the Formal Sciences II: Applications of Mathematical Logic in Philosophy and Linguistics*, pages 163–174. Kluwer, Dordrecht, 2003.
 - [46] Juliano Maranhão. Refinement. A tool to deal with inconsistencies. In *Proc. of the 8th ICAIL*, pages 52–59, 2001.
 - [47] Juliano Maranhao. Some Operators for Refinement of Normative Systems. In *Proceedings of Jurix 2001*, pages 103–115. IOS Press, 2001.
 - [48] Juliano Maranhão. Why was Alchourrón afraid of snakes? *Análisis Filosófico*, XXVI(1):62–92, 2006.
 - [49] Juliano Maranhão. Conservative Contraction. In *The many sides of logic*, pages 465–479. College Publications, 2009.
 - [50] Juliano Maranhão. A logical architecture for dynamic legal interpretation. In *Proceedings of the Eight International Conference on AI and Law ICAIL '17*,, pages 129–38. ACM Press, 2017.
 - [51] Juliano Maranhao. *Positivismo jurídico lógico-incluyente*. Marcial Pons, 2017.
 - [52] Juliano Maranhão and Edelcio G. de Souza. Contraction of combined normative sets. In *Deontic Logic and Normative Systems: 14th International Conference, DEON 2018*, pages 247–261. Springer, 2018.
 - [53] Juliano Maranhão and Giovanni Sartor. Value assessment and revision in legal interpretation. In *Proceedings of the 17th International Conference on Artificial Intelligence and*

- Law, ICAIL 2019*, pages 219–223, New York, New York, USA, jun 2019. Association for Computing Machinery, Inc.
- [54] Andrei Marmor. *Interpretation and Legal theory*. Hart, 2005.
 - [55] Pablo E Navarro and Jorge L Rodriguez. *Deontic Logic and Legal Systems*. Cambridge University Press, 2014.
 - [56] Xavier Parent. Moral particularism in the light of deontic logic. *Artificial Intelligence and Law*, 19:75–98, 2011.
 - [57] Xavier Parent and Leendert van der Torre. *Input/output logics*. College Publications, London, 2013.
 - [58] Xavier Parent and Leendert van der Torre. The pragmatic oddity in norm-based deontic logics. In *Proceedings of the International Conference on Artificial Intelligence and Law*, pages 169–178. Association for Computing Machinery, jun 2017.
 - [59] Xavier Parent and Leendert van der Torre. - input/output logics with a consistency check. In *Proceedings of the 14th International Conference on Deontic Logic and Normative Systems (DEON2018)*, 2018.
 - [60] Aleksander Peczenik. *On Law and Reason*. Kluwer, 1989.
 - [61] Henry Prakken, Adam Wyner, Trevor R. Bench-Capon, and Katie Atkinson. A formalisation of argumentation schemes for legal case-based reasoning in ASPIC+. *Journal of Logic and Computation*, 25:1141–1166, 2015.
 - [62] Joseph Raz. Legal Positivism and the Sources of Law. In *The Authority of Law*, pages 37–52. Oxford University Press, 1979.
 - [63] Joseph Raz. Authority, law and morality. *The Monist*, 68(3):295–324, jul 1985.
 - [64] Joseph Raz. *Authority and Interpretation*. Oxford University Press, 2009.
 - [65] Jandson S. Ribeiro, Abhaya Nayak, and Renata Wassermann. Towards belief contraction without compactness. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, pages 287–296, 2018.
 - [66] Antonino Rotolo. *Identità e somiglianza: saggio sul pensiero analogico nel diritto*. Clueb, 2001.
 - [67] Hans Rott. Two dogmas of belief revision. *Journal of Philosophy*, 97(9):503–522, 2000.
 - [68] Giovanni Sartor. The Nature of Legal Concepts: Inferential Nodes or Ontological Categories. In Gianmaria Ajani, Ginevra Peruginelli, Giovanni Sartor, and Daniela Tiscornia, editors, *Proceeding of the Conference on “Approaching the Multilanguage Complexity of European Law: Methodologies in Comparison”*. European Press Academic Publishing, 2007.
 - [69] Giovanni Sartor. The logic of proportionality: Reasoning with non-numerical magnitudes. *German Law Journal*, pages 1419–57, 2013.
 - [70] Giovanni Sartor. Consistency in balancing: from value assessments to factor-based rules. In D. Duarte and S. Sampaio, editors, *Proportionality in Law: An Analytical Perspective*, pages 121–36. Springer, 2018.
 - [71] John R Searle. *The Construction of Social Reality*. Free Press, 1995.

- [72] Lawrence B. Solum. The interpretation-construction distinction. *Constitutional Commentary*, 27:95–118, 2010.
- [73] Audun Stolpe. Norm-system revision: theory and application. *Artificial Intelligence and Law*, 18:247–283, 2010.
- [74] Luciano H. Tamargo, Diego C. Martinez, Antonino Rotolo, and Guido Governatori. Temporalised belief revision in the law. In *Frontiers in Artificial Intelligence and Applications*, volume 302, pages 49–58. IOS Press, 2017.
- [75] Luciano H. Tamargo, Diego C. Martínez, Antonino Rotolo, and Guido Governatori. Time, defeasible logic and belief revision: Pathways to legal dynamics. *FLAP*, 8(4):993–1022, 2021.
- [76] Hans Van Ditmarsch and Barteld Kooi. The secret of my success. *Synthese*, 151(2):201–232, Jul 2006.
- [77] Jerzy Wróblewski. Legal Language and Legal Interpretation. *Law and Philosophy*, 4:239–255, 1985.
- [78] Jerzy Wróblewski. *The Judicial Application of Law*. Kluwer, 1992.
- [79] Zhiqiang Zhuang, James P. Delgrande, Abhaya C. Nayak, and Abdul Sattar. Reconsidering AGM-style belief revision in the context of logic programs. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, pages 671–679, 2016.
- [80] Zhiqiang Zhuang, Zhe Wang, Kewen Wang, and Guilin Qi. DL-lite contraction and revision. *J. Artif. Intell. Res.*, 56:329–378, 2016.