Legal Knowledge and Information Systems
R. Markovich et al. (Eds.)

doi:10.3233/FAIA251577

Contrary-to-Duty Rights: From Hohfeld to Agreement Revision

Huimin DONG [a], Leendert VAN DER TORRE [b,c] and Liuwen YU [b]
[a] *TU Wien, Vienna, Austria*
[b] *University of Luxembourg, Esch-sur-Alzette, Luxembourg*
[c] *Zhejiang University, Hangzhou, China*

Abstract. We present a rights-first model of contrary-to-duty (CTD) reasoning with two remedial regimes and a revision track. In the *CTD-Claim* regime, when a primary duty is not fulfilled and no exception applies, a remedial claim detaches *automatically*, without any recognition act. In the *CTD-Power* regime, a remedial claim arises only if the rights-holder *exercises* a recognition power; until then there is no recognised violation and no remedial duty. *Revision-of-Duty* (RoD) is an alternative discretionary power that adapts the primary duty without recognising a violation, keeping the purpose aligned and avoiding sanctions. Under CTD-Power, *revision* competes directly with *recognition* on the same case. We develop a model for agreements, give concise dynamic-logic-style specifications of guards and acts, implement an institutional rights system that executes these specifications over live Hohfeldian bundles (with per-case exclusivity, exception handling and provenance), and show how agentic AI can use reasons to choose among the admissible acts within the rights-first framework.

Keywords. Hohfeldian Rights; Contrary-to-Duty reasoning; CTD-Claim; CTD-Power; Revision-of-Duty; Agreement Technologies; Normative Multi-Agent Systems

1. Introduction to the Rights-First Perspective

Contrary-to-duty (CTD) reasoning is usually presented in the context of *obligations*: if a primary obligation is not met, a secondary one applies [1, 2, 3, 4, 5]. Among Hohfeld's four correlatives of rights — *claim/duty*, *privilege (liberty)/no-claim*, *power/liability*, and *immunity/no-power* [6] — obligations correlate with only one correlative in legal settings, which is *claim/duty*. We study *contrary-to-duty (CTD) rights* as directed legal positions that become operative *because* some primary duty φ was not fulfilled ($\neg\varphi$).

For example, let a lease require the tenant to pay rent on the first day of each month (φ). On the fifth day, payment has still not been made ($\neg\varphi$). This familiar situation are supported by the standard CTD forms:

- **CTD obligation (duty-first baseline).** "If rent is not paid on the first day of the month, the tenant must pay a late fee."
- **CTD permission (duty-first variant).** "If the landlord fails to provide heating, the tenant may withhold rent."

Now read the same situation through Hohfeld's four corners as *CTD-rights* that come into play *because* non-fulfillment $\neg\varphi$ occurs:

- **CTD-*Claim*.** On $\neg\varphi$ (no payment by the first day), and in the absence of exceptions, a *remedial claim* (e.g. late fee) *detaches automatically*; the landlord holds a claim, the tenant a correlative duty.
- **CTD-*Privilege* (Liberty).** On $\neg\varphi_L$ (e.g. no heating), the tenant holds a *privilege* to withhold rent; correlatively the landlord has *no-claim* to payment during the outage (a CTD permission as a Hohfeld legal relation).
- **CTD-*Power*.** On $\neg\varphi$, the lease gives the landlord a *recognition/enforcement power*: a remedial claim *arises only if* that power is exercised (there is no automatic duty otherwise).
- **CTD-*Immunity*.** On $\neg\varphi$ with a recognised *exception* (e.g. force majeure concerning payroll), the tenant holds an *immunity*: the landlord's recognition power is *disabled* in this case (no sanction can be validly created).

This example suggests that the traditional CTD obligation and permission are just two facets of a richer family of *CTD-rights*, and that their behaviours differ precisely along Hohfeld's correlatives.

Here, we study whether this plurality matters for agreements [7] in normative multi-agent systems (MAS) [8]. In our study, we observe that, after distinguishing the four CTD-rights, a crucial form of *power* comes into focus: *revision-of-duty* (RoD).

- **Revision-of-Duty (RoD)**. The parties *revise* the rent date (e.g. move it to the tenth day of this month). The purpose of the lease is preserved; the tenant is *immune* from violation-recognition with regard to the original date.

Figure 1 depicts our focus. CTD-CLAIM is situated at the *claim/duty* corner (a remedial claim detaches on $\neg\varphi$ unless an exception holds). CTD-POWER is situated at the *power/liability* corner (a remedial claim appears only when a recognition power is exercised). *RoD* is a *revision* power that changes the primary duty so that the same $\neg\varphi$ no longer counts as a violation and an *immunity* against recognition is created.

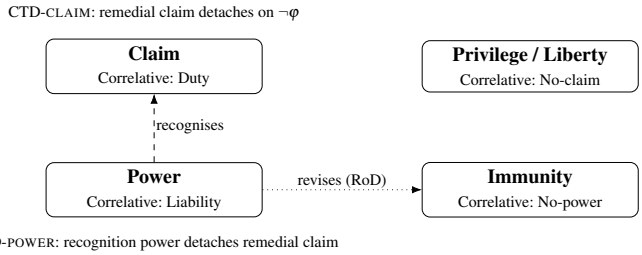

Figure 1. Hohfeldian rights with CTD markers. CTD-CLAIM: remedial claim detaches on non-fulfilment. CTD-POWER: remedial claim appears only if a recognition power is exercised. RoD is a *revision* power that adapts the duty and creates an immunity against recognition of the previous case.

The layout of this paper is as follows. Section 2 integrates our lifecycle intuition by applying CTD-CLAIM, CTD-POWER, and RoD to an agreement example. Section 3 explains how these ideas map to agreement states in an institution (purely at the level of concepts such as claim, power, and immunity). Section 4 briefly discusses what the conceptual picture implies for agentic AI & Law. Section 5 reviews related work, and Section 6 concludes the paper.

2. Specification of Agreements and Hohfedian Rights

We treat a signed agreement as a structured bundle of rights where the correlation between claim and duty is straightforward. To give another example, once the primary rights are activated, a buyer's claim to delivery induces a seller's duty to deliver, and breaches may trigger sanctions. Agreements are not static. Parties can exercise powers to either remedy a breach or revise the agreement in response to an unexpected situation by, for instance, making a delay permissible.

We adopt a *rights-first* perspective, reasoning about the CTD regime of secondaty rights in two distinct ways for the purpose of remedying issues in agreements. Building on the Hohfeldian rights with CTD markers in Figure 1, Figure 2 summarises the lifecycle we adopt of rights-based agreements: from *creation* to *primaries active*, an *issue* may trigger either the CTD track (with its two regimes) or the RoD track. A *revision* creates a new bundle and loops back to the active-bundle state, allowing *multiple revisions over time*. [1] This mirrors how agreements in practice may undergo several targeted amendments before final discharge, while keeping provenance auditable and conflicts controlled. The overall setting aligns with Agreement Technologies [7] and normative multi-agent systems (e.g. [8]), but uses Hohfeldian rights not merely for analysis but also to *operate* the state and evolution of agreements.

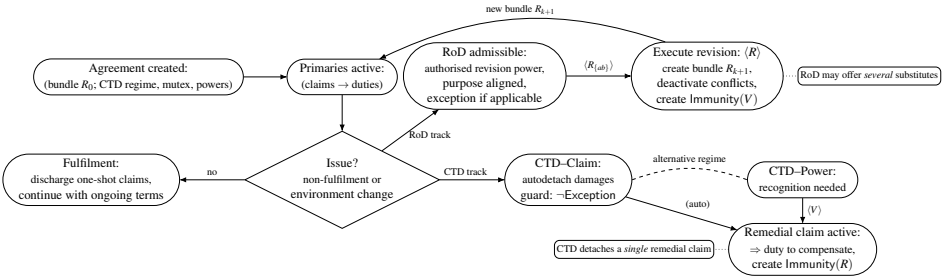

Figure 2. The agreement lifecycle with feedback. Its creation fixes issues with the regime and relevant powers. During execution, an issue triggers either CTD (automatic claim or recognition-based) or RoD (revision). CTD creates an immunity against revision; RoD creates an immunity against recognition and loops back to the *primaries active* state with the new bundle R_{k+1}, thereby allowing multiple revisions. Mutual-exclusion guards prevent conflicting content from co-existing; discharge/persistence govern how claims end.

CTD-Claim Regime: Automatic Detachment. In this first CTD-regime, secondary rights arise automatically. Suppose the seller fails to deliver on time without exception: the agreement stipulates that the buyer immediately holds a remedial claim (e.g. compensation). This is the CTD-Claim path. Under Hohfeld's framework, a remedial claim for the buyer entails a correlative duty for the seller. Importantly, no discretionary step is involved; the secondary duty detaches by operation of the law.

CTD-Power Regime: Recognition-Based Detachment. The second CTD-regime introduces discretion. Here, no remedial duty exists unless the rights-holder actively recognises a violation. In this CTD-Power regime, the buyer possesses a power to enforce, but until *exercised*, no claim arises. Recognition thus functions as a gatekeeping act: it transforms non-fulfilment into a violation and detaches the corresponding remedial duty.

[1] CTDs can repeat at multiple levels and may involve several layers of violation; see the detailed review in [9] for further discussion.

Rights (Hohfeld bundles)	Claim	Power	Recognition	Revision	Programs (Power execution)
	$\mathsf{Claim}_{a\leftarrow b}(p/\neg q)$	$\mathsf{Power}_{a\leftarrow b}(p/\neg q)$	$\langle \mathsf{V}_b \rangle$	$\langle \mathsf{R}_{\{ab\}} \rangle$	
Language $\mathscr{L}$ is defined as $\varphi := p \in \mathit{Prop} \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \mathsf{Claim}_{a\leftarrow b}(\varphi/\varphi) \mid \mathsf{Power}_{a\leftarrow b}(\varphi/\varphi) \mid \langle \mathsf{V}_b \rangle \varphi \mid \langle \mathsf{R}_{\{ab\}} \rangle \varphi$					

Table 1. Representation of Rights-First System = Hohfeld bundles Rights + Power Programs.

Revision-of-Duty Track: Purpose-Preserving Adaptation. Revision-of-Duty (RoD) is an alternative power that allows parties to adapt the agreement itself. Instead of treating non-fulfillment as a violation, the situation may be re-framed as an exception, prompting joint *revision*. For example, if delayed payment results from a temporary budget freeze, the parties may agree to extend the deadline. No breach is recognised, no sanction imposed, and the contract continues with a revised bundle of rights. RoD therefore embodies a *discretionary*, forward-looking power: it preserves contractual purpose and *immunises* the case from violation-recognition.

RoD is significant because it reveals a dimension absent from standard CTD analyses: *discretionary power*[2] to adapt duties. This distinguishes remedial enforcement from cooperative adjustment, and it highlights that the rights-first perspective is richer than duty-first models. Our specification is a *conceptual* contribution — it will discuss this conceptual distinction through a set of logical constraints while remaining intentionally minimal and without fully embedding formal semantics.

Seen through a Hohfeldian *rights-first* lens, this reasoning system provides a way to track whether situations are treated as *violations* or as *exceptions*. Accordingly, *agents* must decide whether to exercise their legal powers by recognising a violation or adopt a revision to cope with an exception.

The core elements of specifying rights-first reasoning are twofold: (i) rights as conditional norms [11], and (ii) programs that represent the exercise of those powers — such as violation recognition and agreement revision — as summarised in Table 1. Following Hohfeld's distinction between static and dynamic rights, we categorise rights into two main groups: claims and powers. Other legal rights such as privilege and immunity can be defined from these categories through Hohfeldian correlative relations.[3] The representation of programs extends this framework by capturing the idea of *exercising power*, thereby foregrounding the role of *agents' choices* within a dynamic logic formalisation [12].[4] For instance, $\langle \mathsf{V}_A \rangle \mathsf{Claim}_{M\leftarrow A}(\mathsf{Compensation}/\mathsf{NF})$ means that agent A's exercise of violation-recognition results in a claim from A against M for compensation when the primary right is not fulfilled. By contrast, $\langle \mathsf{R}_{\{AM\}} \rangle \neg \mathsf{Claim}_{M\leftarrow A}(\mathsf{Compensation}/\mathsf{NF})$ means that agents A and M's joint revision leads to the negation of this claim-right. In this work, we focus on the specification; the semantics will be addressed in future work.

[2] Hart [10] maintained that the concept of discretion plays an important role in the legal domain. It involves making a choice that remains bound by legal values and principles, while keeping alternative options possible.

[3] Given the language $\mathscr{L}$, the propositional connectives $\vee, \rightarrow, \leftrightarrow$ can be defined in a standard way. By Hohfeld's system of correlatives, other fundamental legal positions can be defined as follows: $\mathsf{Duty}_{a\leftarrow b}(\varphi/\psi) := \mathsf{Claim}_{b\leftarrow a}(\varphi/\psi)$, $\mathsf{Liberty}_{a\leftarrow b}(\varphi/\psi) := \neg\mathsf{Claim}_{b\leftarrow a}(\neg\varphi/\psi)$, $\mathsf{Liability}_{a\leftarrow b}(\varphi/\psi) := \mathsf{Power}_{b\leftarrow a}(\varphi/\psi)$, and $\mathsf{Immunity}_{a\leftarrow b}(\varphi/\psi) := \neg\mathsf{Power}_{b\leftarrow a}(\varphi/\psi)$. For convenience, we write $\mathsf{Right}_{a\leftarrow b}(\varphi/\psi)$ to denote any element of this Hohfeldian bundle. The program operator $[\cdot]$ is the dual of $\langle\cdot\rangle$, e.g. $[\mathsf{R}_{\{ab\}}]\varphi := \neg\langle \mathsf{R}_{\{ab\}} \rangle \neg\varphi$.

[4] Here, we adopt the distinction between *legal competence* and *ontic ability* proposed by Markovich [13], and we focus on a lightweight representation of power programs as the exercise of legal competence. In this way, our specification remains as general as possible while still being able to capture the distinctions highlighted in existing frameworks.

Example (Public Park Contract). Suppose Municipality M contracts Contractor A to renovate a public park (Table 1).

Primary right: A has the right to payment on time: $\mathsf{Right}_{M\leftarrow A}(\mathsf{Pay\text{-}timely}/\top)$.

If M fails to pay on time (NF), two kinds of secondary rights arise: claims or powers.

CTD-claim path: Contractor A may automatically obtain a claim-right against M for compensation given the non-fulfillment — that is $\mathsf{Claim}_{M\leftarrow A}(\mathsf{Compensation}/\mathsf{NF})$.

Alternatively, the agreement may be structured so that Contractor A holds a power — $\mathsf{Power}_{M\leftarrow A}(\mathsf{Compensation}/\mathsf{NF})$. This pathway leads to two different *agentic* tracks:

CTD-power path: A has no immediate claim-right to compensation. The right materialises only if A exercises its violation-recognition power, which then leads to a claim, as in the CTD-claim path: $\langle \mathsf{V}_A\rangle \mathsf{Claim}_{M\leftarrow A}(\mathsf{Compensation}/\mathsf{NF})$.

RoD-power path: Alternatively, M and A may classify the situation as an exception (EXP), e.g. they may extend the deadline because of a budgetary freeze. In this case, no violation is recognised, compensation is not triggered, and the contract adapts to preserve its purpose. *Exercising this revision power* [5] yields a new right: $\langle \mathsf{R}_{\{AM\}}\rangle \mathsf{Claim}_{M\leftarrow A}(\mathsf{30\text{-}days}/\mathsf{EXP})$.

To address conflicts, we introduce a lightweight representation of *consistency* principles for reasoning about Hohfeldian bundles, inspired by Hohfeld's (1913) theory [6].

- No contradictory claim-rights: (C_1) $\neg(\mathsf{Claim}_{a\leftarrow b}(\varphi/\psi) \wedge \mathsf{Claim}_{a\leftarrow c}(\neg\varphi/\psi))$;
- Immunity as absence of power: (C_2) $\langle \mathsf{V}_b\rangle[\mathsf{R}_{\{ab\}}]\bot$ and $\langle \mathsf{R}_{\{ab\}}\rangle[\mathsf{V}_b]\bot$.

Principle (C_1) rules out conflict scenarios in which two contractors simultaneously demand mutually exclusive compensation from the Municipality when only one official contractor exists. Principle (C_2) refines the notion of immunity in terms of the incompatibility of alternative actions: if a violation-recognition is acted upon, then no revision is executable, and vice versa — a form of *CTD–RoD program exclusivity*.

We impose three constraints that prevent conflicts among CTD-claims, CTD-powers, and RoD-powers and that ground our operational framework. (P_1) *Auto-remedy excludes contrary powers*: if a CTD-Claim for φ is active (automatic detachment), then no power is available that would, if exercised, create a claim for an incompatible ψ. (P_2) *Mutually contrary CTD-powers do not co-exist*: two CTD-Powers whose exercise would yield mutually incompatible remedial claims must not be simultaneously available for the same case and locus (i.e. the same right-holder and counterparty). (P_3) *Revision power may co-exist with a contrary CTD-power* (choice point): a RoD-power to create φ may co-exist with a CTD- or RoD-power whose exercise would create a mutually incompatible ψ. Co-existence is allowed only at the level of availability; exercising one disables the other via track-specific immunities. Thus P_1 and P_2 enforce consistency, while (P_3) preserves the discretionary choice that is central to the rights-first view.

We briefly comment on the specification of *exception* (EXP), the central notion in the rights system. If we wish to explicitly formalise the distinction between *exception* and *violation* — where exceptions (i) exclude sanctions and (ii) block recognition — two illustrative implications are: (1) $\mathsf{EXP} \rightarrow \langle \mathsf{R}_{\{ab\}}\rangle \neg \mathsf{Claim}_{a\leftarrow b}(\mathsf{sanction}/\psi)$ and (2) $\mathsf{EXP} \rightarrow \neg\langle \mathsf{V}_b\rangle\top$. That is, exceptions activate revision without sanction but never recognition.

[5] A revision power may be exercised individually or jointly, thereby detaching a new bundle of Hohfeldian rights. In contrast to existing frameworks [5, 14, 15], our lightweight specification provides a novel representation of rights detachment in multi-agent settings,

3. Institutional Rights System for Managing Agreements

The rights system we introduce in this section is a proof-of-concept institutional layer that maps observed *facts* and recorded *legal acts* to a live Hohfeld bundle for each agreement, while satisfying the constraints of the previous section.

Inputs are facts (performance logged, deadline passed), acts (RECOGNISE breach; authorised REVISE; ANNUL contract), and creation-time configuration (primary contents, CTD regime, admissible revisions, mutex constraints). Outputs are current **claims** (and induced **duties**), available **powers**, and any **immunities**, each with explicit *provenance* (*primary*, *CTD-claim*, *CTD-power+recognised*, *revision*). The system never chooses between admissible acts — that responsibility rests with the parties. By design, there is no duty without a live claim, every claim has schema or power provenance, CTD–RoD exclusivity applies on a per-case basis (via immunities), and exceptions disable recognition and create no sanctions. This stays faithful to the rights-first approach: *claims* and *powers* are primitive; *duties*, *liabilities*, and *immunities* arise as correlatives in states.

A useful analytical strategy is to separate *facts* from *acts*. Facts move the lifecycle without exercising power: *performance* discharges a one-shot normative element (e.g. primary, remedy or substitute); an *issue* (non-fulfilment) opens a case; an *exception* blocks recognition. Acts exercise powers and thus change bundles: RECOGNISE detaches the remedy (when available) and creates $\mathsf{Imm}(R)$; joint REVISE creates a substitute bundle and $\mathsf{Imm}(V)$; WAIT leaves the bundle unchanged; and ANNUL terminates the contract, leaving no rights remaining with the parties. This cleanly implements the conceptual contrast in Section 2 between *automatic* CTD-CLAIM and *discretionary* CTD-POWER.

Table 2 gives the state semantics.

- **S0 (Primaries)**: preconditions hold; primary claims induce duties; if configured, a revision power is available.
- **S1 (Discharge)**: a performed one-shot element — primary, remedy, or substitute — is removed; persistent terms continue.
- **S2 (Case open)**: non-fulfilment is recorded and the CTD regime is evaluated. In **CTD-Claim**, unless an exception holds, the remedy *detaches automatically* to **S3a**; there is no agent choice. In **CTD-Power**, the remedial posture depends on the *exercise* of that power; a **Choice** appears (**S3b**).
- **S3a (Remedy auto)**: remedy live and $\mathsf{Imm}(R)$ block revision on the case.
- **S3b (Choice)**: RECOGNISE (to **S4a**), joint REVISE (to **S4b**), or WAIT (self-loop).
- **S4a (Recognised)**: remedy active + $\mathsf{Imm}(R)$. In **S1**, **S4a** is discharged via settlement.
- **S4b (Revised)**: substitute bundle created; conflicts deactivated under mutex; $\mathsf{Imm}(V)$ on the old case; loop back to **S0** so multiple, purpose-preserving revisions are possible.
- **S5 (Exception)**: $\mathsf{Imm}(V)$; recognition disabled; consensual revision still possible to preserve purpose, or the contract can be terminated.

Table 3 names the transitions and separates guards from acts. Fact-transitions are *performance* (S0→S1), *issue* (S0→S2), *exception* (S2→S5). Act-transitions are RECOGNISE (S3b→S4a), joint REVISE (S3b→S4b; S5→S4b), WAIT (S3b→S3b), *settlement* (S3a/S4a→S1), and *termination* (S5→S1). A configuration flag sets remedies as *alternative* (the primary is discharged once a remedy is activated) or *cumulative* (the primary remains until the issue is resolved); mutex guards prevent contradictory duties. Exceptions differ from licences: a true exception creates $\mathsf{Imm}(V)$ and excludes sanctions; a mere licence may suspend a primary but, when it lapses, recognition is not disabled by that licence alone. RoD transitions depend on authorisation (who may exercise the re-

ID	Name	What it means	Bundle highlights
S0	Primaries active	Preconditions met; no mutex conflict.	Claims induce duties; revision power available (if configured).
S1	Fulfilled / discharge	A one-shot element is performed.	Discharge that particular claim/duty; next period/case.
S2	Case open (non-fulfilment)	Due date passes or performance fails.	CTD-Claim: auto remedy; CTD-Power: choice emerges.
S3a	Remedy active (auto)	Auto CTD-claim detached.	Remedy live + $\mathsf{Imm}(R)$ (blocks revision).
S3b	Choice (CTD-Power)	No exception; CTD-Power regime.	RECOGNISE / REVISE (joint) / WAIT.
S4a	Recognised	Recognition exercised.	Remedy live + $\mathsf{Imm}(R)$.
S4b	Revised	Joint revision created.	New primaries; $\mathsf{Imm}(V)$; loop to S0.
S5	Exception	True exception.	$\mathsf{Imm}(V)$; no sanctions.

Table 2. Rights system states. Each state is a *bundle condition* indicating which claims, powers, and immunities exist now and why.

From → To	Trigger (guard ∧ act)	Effect on bundle	Notes
S0 → S1	Performance	Discharge performed claim/duty.	Next period/case → S0.
S0 → S2	Issue	Open a case.	Regime/exception evaluated.
S2 → S3a	CTD-Claim ∧¬Exc	Activate remedy + $\mathsf{Imm}(R)$.	No agent choice.
S2 → S3b	CTD-Power ∧¬Exc	Offer recognise / revise / wait.	Choice point.
S2 → S5	Exception	Create $\mathsf{Imm}(V)$.	Recognition disabled.
S3a/S4a → S1	Remedy settled	Discharge remedy.	Case closes.
S3b → S3b	WAIT	Bundle unchanged.	Self-loop.
S3b → S4a	RECOGNISE	Remedy + $\mathsf{Imm}(R)$.	Blocks revision on this case.
S3b/S5 → S4b	REVISE (as authorised)	Substitute bundle + $\mathsf{Imm}(V)$.	Loop to S0.
S4b → S1	Substitute fulfilled	Discharge substitute.	Then S0.
S5 → S1	Purpose failed	Terminate the contract.	Then S0.

Table 3. Named transitions. Facts move the case; acts exercise powers. Immunities enforce per-case CTD/RoD exclusivity and document why a path is closed.

vision power) rather than on exception per se. Exceptions explain why sanction is not appropriate and so justify revision, but the power to revise is conferred by the agreement.

Figure 3 is read in three threads aligned with Section 2. *CTD-Claim* is automatic: from **S2**, the remedy becomes active (**S3a**) and revision is blocked by $\mathsf{Imm}(R)$. Settlement leads to **S1**, then the "next period/case" edge returns to **S0**, representing multiple CTD episodes *across* cycles. *CTD-Power* is discretionary: **S3b** presents a real choice — RECOGNISE to **S4a**, joint REVISE to **S4b**, or WAIT. *RoD* is joint and forward-looking: **S4b** loops to **S0** with a new bundle, which enables targeted amendments over time while provenance and mutex constraints remain auditable. Case identity is clear: each due item (or declared issue) opens exactly one case; exclusivity is case-specific, and each case cycles naturally across periods. Concurrency is handled institutionally: the first committed act is protected and subsequent conflicting acts are rendered inadmissible.

In sum, the rights system operationalises Sections 1–2. CTD-*Claim* involves no choice. CTD-*Power* creates the discretionary space where RECOGNISE, REVISE, and WAIT genuinely compete on the same case. *RoD* resets the bundle while preserving the old track. The result is a short, auditable path from facts and recorded acts to live rights and duties, with invariants and exclusivity guaranteed by design.

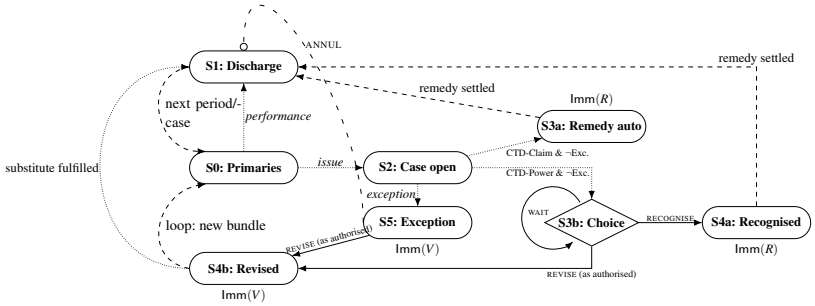

Figure 3. The rights system in a compact state machine. CTD-*Claim* offers no choice; CTD-*Power* offers a choice. *RoD* loops to **S0** (with a new bundle). $\mathsf{Imm}(R)$ blocks revision on the case; $\mathsf{Imm}(V)$ blocks recognition. Dotted lines denote facts, dashed lines the configuration, and solid lines legal acts.

4. Collaborative Agentic AI System for Rights-Based Agreements

Section 3 showed that a rights-first model suffices for the *institutional* system: the rights system computes, from facts and recorded acts, which claims and powers are admissible at each stage. We now show that the same approach suffices for an *agent* system concerned with practical reasoning. Taken together, the institutional and agent systems form a multi-agent system: the rights system provides admissible options while the agent system selects one of them for execution, both governed by the logical constraints in the specification. This is our second proof of concept: an *agentic AI in Law* can operate within the rights-first framework by adding a transparent, reason-based choice mechanism for exercising discretionary powers and producing a single executable act.

We use *agentic AI* in the general sense of systems able to set goals, plan, call external tools, coordinate with others, and maintain working memory, often implemented with Large Language Model (LLM) components [16]. In our setting, an agentic AI is an *autonomous program* representing a contractual party. It interfaces with the rights system and exercises discretion by selecting from admissible acts. LLM components may assist as personal assistants, e.g. retrieving contractual grounds, organising evidence, or drafting auditable explanations, but the assignment of normative weights remains a matter of policy, precedent, and evidence.

To explain how an agent selects from admissible acts, we adopt Tucker's *dual-scale theory of reasons* [17]. A single-scale picture collapses all reasons into one number and presumes a fixed exchange between "counts-in-favor" and "makes-obligatory", which distorts normative phenomena. Dual scale separates these forces. Each agent identifies relevant grounds (liquidity, fairness, reputation) and activates them contrastively: a ground becomes a reason for a only in relation to an alternative b, and its weight may vary across comparisons (*contrastivism* [17, p. 122]). Reasons come with three weight values: (i) *justifying weight (JW)*, the degree to which an act is pushed towards permissibility; (ii) *requiring weight (RW)*, the degree to which an act is required, pushing alternatives towards impermissibility; and (iii) *commending weight (CW)*, the degree to which one permissible act is morally preferable to another: "altruistic benefits always commend more than self-interested benefits of the same size" [17, p. 172]. The Permission Scale checks whether one JW outweighs the RW of its rivals, while the Commitment Scale mirrors it by flipping sides. A requirement arises when an act is permissible and all alternatives are impermissible.

Formal rules. For an agreement ι, $\mathscr{A}_\iota = \{\textsc{recognise}, \textsc{wait}\} \cup \{\textsc{revise}(r) \mid r \in R_\iota\}$ denotes the set of admissible acts. For any ordered pair (a,b) of distinct acts, let $\mathscr{G}_\iota(a \parallel b)$ be the contrast-indexed set of grounds that support act a against act b. Each $g \in \mathscr{G}_\iota(a \parallel b)$ contributes one or more reasons with weight $w_g \in \mathbb{N}$ ranging over natural numbers:

$$\mathsf{JW}(a \parallel b) = \sum_{g \in \mathscr{G}_\iota^{\text{just}}(a \parallel b)} w_g, \quad \mathsf{RW}(a \parallel b) = \sum_{g \in \mathscr{G}_\iota^{\text{req}}(a \parallel b)} w_g, \quad \mathsf{CW}(a) = \sum_{g \in \mathscr{G}_\iota^{\text{com}}(a)} w_g.$$

Dual-scale verdicts are pairwise: a is *permissible* against b iff $\mathsf{JW}(a \parallel b) \geq \mathsf{RW}(b \parallel a)$, has a *commitment* against b iff $\mathsf{RW}(a \parallel b) > \mathsf{JW}(b \parallel a)$, and is *required* if both hold. With multiple options, permissibility and requirement mean winning all pairwise contests. Commending weight is not used to determine the deontic status of acts but only for ranking within the permissible set, thus refining choice (akin to advice or soft-law guidance). If no act is permissible, the default is WAIT.

Running example (Public Park Contract). Contractor A delivered, M missed timely payment. Admissible acts: RECOGNISE (REC), WAIT, REVISE(30d) (REV30D), and REVISE(14d) (REV14D). We show three contrasts. Each row lists one active and contrast-indexed reason for the *left* option with its (jw, rw). In Contrast 3, where we have two tied permissible acts, commending weight comes in to make the decision.

Contrast 1: REVISE(30d) vs. RECOGNISE

$(g, a \,\|\| b)$	Justifying weight	Requiring weight	Comment
(purpose, REV30D $\|\|$ REC)	2	1	preserve purpose; avoid closure
(liquidity$_A$, REV30D $\|\|$ REC)	2	1	utility/efficiency for A
(reputation$_M$, REV30D $\|\|$ REC)	1	2	sanction harms M, reputation/compliance
(compliance, REC $\|\|$ REV30D)	1	3	policy duty
(liquidity$_A$, REC $\|\|$ REV30D)	2	0	immediate liquidity

Totals: JW(REV30D $\|$ REC) $= 5 \geq$ RW(REC $\|$ REV30D) $= 3$, so REV30D is permissible; RW(REV30D $\|$ REC) $= 4 \geq$ JW(REC $\|$ REV30D) $= 3$, so REV30D has a commitment. Thus, rev30d is required. Meanwhile, JW(REC $\|$ REV30D) $= 3 <$ RW(REV30D $\|$ REC) $= 4$, so REC is not permissible; and RW(REC $\|$ REV30D) $= 3 <$ JW(REV30D $\|$ REC) $= 5$, so REC is not a commitment. Thus, REC is not required.

Contrast 2: REVISE(30d) vs. WAIT

$(g, a \,\|\| b)$	Justifying weight	Requiring weight	Comment
(opCont, REV30D $\|\|$ WAIT)	2	1	keep services, utility
(budget, REV30D $\|\|$ WAIT)	1	1	administrative timing
(uncertainty, WAIT $\|\|$ REV30D)	1	0	information-gathering, uncertainty

Totals: JW(REV30D $\|$ WAIT) $= 3 \geq$ RW(WAIT $\|$ REV30D) $= 0$, so REV30D is permissible; RW(REV30D $\|$ WAIT) $= 2 \geq$ JW(WAIT $\|$ REV30D) $= 1$, so REV30D has a commitment. Thus, REV30D is required. Meanwhile, JW(WAIT $\|$ REV30D) $= 1 <$ RW(REV30D $\|$ WAIT) $= 2$, so WAIT is not permissible; and RW(WAIT $\|$ REV30D) $= 0 <$ JW(REV30D $\|$ WAIT) $= 3$, so no commitment holds for WAIT. Thus, WAIT is not required.

Contrast 3: REVISE(14d) vs. REVISE(30d)

$(g, a \,\|\| b)$	Justifying weight	Requiring weight	Comment
(fairness$_A$, REV14D $\|\|$ REV30D)	3	1	fairness to A, cooperation
(liquidity$_A$, REV14D $\|\|$ REV30D)	2	0	faster cash, utility
(cashflow$_M$, REV30D $\|\|$ REV14D)	1	0	M's short-term stress, utility

Totals: JW(REV14D $\|$ REV30D) $= 5 \geq$ RW(REV30D $\|$ REV14D) $= 0$, so REV14D is permissible; RW(REV14D $\|$ REV30D) $= 1 =$ JW(REV30D $\|$ REV14D) $= 1$, so REV14D is not required. Meanwhile, JW(REV30D $\|$ REV14D) $= 1 \geq$ RW(REV14D $\|$ REV30D) $= 1$, so REV30D is permissible. But RW(REV30D $\|$ REV14D) $= 0 <$ JW(REV14D $\|$ REV30D) $= 5$, so no commitment holds for REV30D. Both are permissible. However, commending weight favours fairness, which counts as an altruistic benefit that tips the choiceworthiness balance.

Given this structure, the agent system can be viewed as a collaborative decision pipeline (Fig. 4) that operationalises the dual-scale model for individual reasoning and shows how multiple agent AIs interact to select a single admissible act.

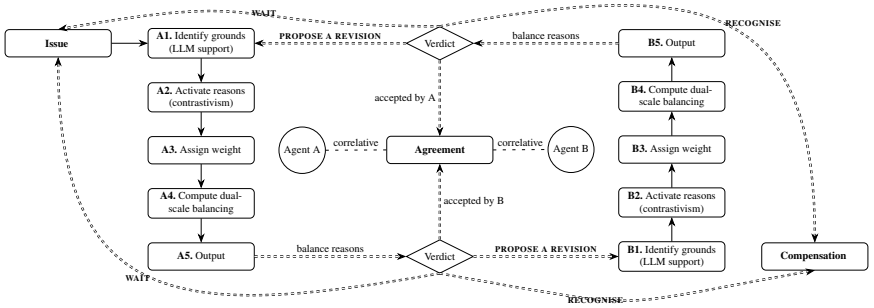

Figure 4. Collaborative agentic AI pipeline for rights-based agreements. Each party is represented by an agent that (1) identifies grounds (with LLM support), (2) activates reasons via contrastivism, (3) assigns weights from policy, precedent, and evidence, (4) computes dual-scale totals, and (5) outputs an act — required, best commended permissible, or WAIT if none are permissible. For verdicts, agents negotiate proposals with mutual acceptance, while recognition and waiting remain unilateral. The diagram shows their revision–accept loop yielding legally intelligible outcomes.

5. Related work

Contrary-to-Duty (CTD) reasoning is central to modelling norm compliance, especially in legal informatics [4, 5]. The traditional *duty-first* view motivates rich logical toolkits — dyadic deontic logic [3, 18], temporal/dynamic deontic logics [19, 20], and non-monotonic approaches (e.g. argumentation [21, 14], constrained I/O logic [22, 23]). Yet, this paradigm only partially reflects the depth of Hohfeld's legal framework [6, 13, 24].

We instead adopt a *rights-first* stance, where claims and powers stand as irreducible primitives. This opens an *operational* perspective on CTD reasoning, akin to system design, which clarifies the *meta*-conditions of compliance — when and which rights are exercised — a dimension missing from duty-first logics.

Our perspective bridges AI and Law by shifting from "*logics about norms*" to "*systems with norms*". Prior work on organisations [25, 19] and agreement technologies [26, 27, 7] model inter-agent complexity via normative multi-agent systems (MAS) [28, 8] and argumentation-based negotiation [29, 30]. Unlike accounts that separate agent-dependence [31] from compliance [32], our specification integrates them via Hohfeld's correlative rights, shifting the focus from *agents* to *rights* and linking the *institutional* (compliance) and *agentic* (trust, reputation, choice) layers at the systems level.

Recent work on *agentic AI* in law addresses tasks from extracting normative statements [33] to supporting knowledge representation [34] and dispute resolution [35] using LLMs. In this light, normative MAS and agreement technologies are not only a conceptual bridge between logics and institutional systems but also a foundation for designing legally intelligible AI agents. Our rights-first approach complements this by assigning LLMs a bounded role — assisting recall, drafting, and explanation — while supporting formal theory and safe integration into legal and contractual systems.

6. Summary

This paper introduces a *rights-first* theoretical framework for agreement lifecycle management that *reconceptualises CTD reasoning*. We take a design-plus-analysis approach: alongside formal study, we build deontic systems that are both executable and auditable. Our focus is on the most persistent challenge — contrary-to-duty (CTD) reasoning. By re-reading CTD reasoning through Hohfeld's lens, we separate (i) *CTD-Claim* (automatic remedial detachment), (ii) *CTD-Power* (recognition-based remedial detachment), and (iii) *Revision-of-Duty* (RoD, purpose-preserving revision), which enables us to design an operational pipeline to implement CTD-based reasoning as an agentic AI system.

Future work includes implementing *rights-based Answer Set Programming (ASP) automation* [15, 36] to align CTD rights with Hohfeldian bundles in practical reasoning [14, 37], and analysing semantics using Deontic Equilibrium Logic [36] to model the modification of rights under formal constraints. We will also study how knowledge and beliefs shape CTD rights, determining when non-fulfilments count as violations, exceptions, or cooperative revisions. A fault-tolerant rights system [28] will ensure robustness under uncertainty while preserving consistency, fairness, and provenance. Another natural step is to move beyond the proof of concept and develop a fuller account of how agentic AI in Law can operate within the rights-first framework: with a specification of how LLMs assist in identifying grounds and reasons, and how agents negotiate using symbolic reasoning and LLMs [38] for revision.

Acknowledgements

We are grateful to Château de Rastignac, where our work first began. We also thank the Luxembourg National Research Fund (FNR) for its support through the project Deontic Logic for Epistemic Rights (DELIGHT) (OPEN O20/14776480). Huimin Dong is supported by the Austrian Science Fund (FWF) and the FNR through the LoDEX project [doi: 10.55776/I6372 and INTER/DFG/23/17415164/LODEX]. Leendert van der Torre acknowledges the financial support of the FNR through the following projects: The Epistemology of AI Systems (EAI) (C22/SC/17111440), DJ4ME – A DJ for Machine Ethics: the Dialogue Jiminy (O24/18989918/DJ4ME), Logical Methods for Deontic Explanations (LoDEx) (INTER/DFG/23/17415164/LoDEx) and Symbolic and Explainable Regulatory AI for Finance Innovation (SERAFIN) (C24/19003061/SERAFIN). Liuwen Yu thanks the FNR for their support through the SERAFIN project (C24/19003061/SERAFIN) and the University of Luxembourg for the Marie Speyer Excellence Grant for the project Formal Analysis of Discretionary Reasoning (MSE-DISCREASON).

References

[1] Chisholm RM. Contrary-to-Duty Imperatives and Deontic Logic. Analysis. 1963;24(2):33-6.

[2] Prakken H, Sergot M. Contrary-to-duty obligations. Studia Logica. 1996;57(1):91-115.

[3] van der Torre L, Tan YH. Contrary-to-duty reasoning with preference-based dyadic obligations. Annals of Mathematics and Artificial Intelligence. 1999;27(1):49-78.

[4] Hage J. Contrary to Duty Obligations A Study in Legal Ontology. In: Proceedings of JURIX 2001; 2001. p. 89-102.

[5] Bench-Capon TJ. Transition systems for designing and reasoning about norms. Artificial Intelligence and Law. 2015;23(4):345-66.

[6] Hohfeld WN. Some Fundamental Legal Conceptions as Applied in Judicial Reasoning. Yale Law Journal. 1913;23(1):16-59.

[7] Ossowski S, editor. Agreement Technologies. Berlin, Heidelberg: Springer; 2013.

[8] Chopra A, van der Torre L, Verhagen H, Villata S, editors. Handbook of Normative Multiagent Systems. College Publications; 2018.

[9] Carmo J, Jones AJ. Deontic logic and contrary-to-duties. In: Handbook of Philosophical Logic: Volume 8. Springer; 2002. p. 265-343.

[10] Hart HL. Discretion. Harvard Law Review. 2013;127(2):652-65.

[11] Gabbay D, Horty J, Parent X, van der Meyden R, van der Torre L, editors. Handbook of Deontic Logic and Normative Systems. College Publications; 2013 (v. 1), 2021 (v. 2).

[12] Harel D, Kozen D, Tiuryn J. Dynamic logic. MIT Press; 2000.

[13] Markovich R. Understanding Hohfeld and formalizing legal rights: the Hohfeldian conceptions and their conditional consequences. Studia Logica. 2020;108(1):129-58.

[14] Atkinson K, Bench-Capon T. Practical reasoning as presumptive argumentation using action based alternating transition systems. Artificial Intelligence. 2007;171(10-15):855-74.

[15] Azzopardi S, Pace GJ, Schapachnik F, Schneider G. Contract automata: an operational view of contracts between interactive parties. Artificial Intelligence and Law. 2016;24(3):203-43.

[16] Acharya DB, Kuppan K, Divya B. Agentic AI: Autonomous intelligence for complex goals–a comprehensive survey. IEEE Access. 2025.

[17] Tucker C. The Weight of Reasons: A Framework for Ethics. New York: Oxford University Press; 2025.

[18] Ciabattoni A, Olivetti N, Parent X. Dyadic obligations: proofs and countermodels via hypersequents. In: Proceedings of PRIMA 2022. Springer; 2022. p. 54-71.

[19] Governatori G, Rotolo A, Sartor G. Temporalised normative positions in defeasible logic. In: Proceedings of ICAIL 2005; 2005. p. 25-34.

[20] Meyer JJC, et al. A different approach to deontic logic: deontic logic viewed as a variant of dynamic logic. Notre Dame J Formal Log. 1988;29(1):109-36.

[21] Bench-Capon TJM. Persuasion in Practical Argument Using Value-Based Argumentation Frameworks. Journal of Logic and Computation. 2003;13(3):429-48.

[22] Makinson D, Van Der Torre L. Constraints for input/output logics. Journal of philosophical logic. 2001;30(2):155-85.

[23] van Berkel K, Markovich R, Straßer C, van der Torre L. Arguing About Choosing a Normative System: Conflict of Laws. In: Sileno G, Spanakis J, van Dijck G, editors. Proceedings of JURIX 2023. vol. 379 of Frontiers in Artificial Intelligence and Applications. IOS Press; 2023. p. 73-82.

[24] Dong H, Roy O. Dynamic logic of legal competences. Journal of Logic, Language and Information. 2021;30(4):701-24.

[25] Jones AJI, Sergot M. A Formal Characterisation of Institutionalised Power. Journal of the IGPL. 1996;4(3):429-45.

[26] Boella G, Pigozzi G, van der Torre L. Five Guidelines for Normative Multiagent Systems. In: Governatori G, editor. Proceedings of JURIX 2009. vol. 205 of Frontiers in Artificial Intelligence and Applications. IOS Press; 2009. p. 21-30.

[27] Boella G, van der Torre L. Reasoning for Agreement Technologies. In: Raedt LD, Bessiere C, Dubois D, Doherty P, Frasconi P, Heintz F, et al., editors. Proceedings of ECAI 2012. vol. 242 of Frontiers in Artificial Intelligence and Applications. IOS Press; 2012. p. 895-6.

[28] Boella G, van der Torre L. Constitutive Norms in the Design of Normative Multiagent Systems. In: Proceedings of CLIMA 2005. Springer; 2005. p. 303-19.

[29] Amgoud L, Kaci S. On the generation of bipolar goals in argumentation-based negotiation. In: Proceedings of ArgMAS 2004. Springer; 2004. p. 192-207.

[30] Liao B, Pardo P, Slavkovik M, van der Torre L. The Jiminy Advisor: Moral Agreements among Stakeholders Based on Norms and Argumentation. Journal of Artificial Intelligence Research. 2023;77:737-92.

[31] Boella G, Caire P, van der Torre L, Villata S. Dependence Networks for Agreement Technologies. In: Ossowski S, Toni F, Vouros GA, editors. Proceedings of AT 2012. vol. 918 of CEUR Workshop Proceedings. CEUR-WS.org; 2012. p. 109-10.

[32] Marosin D, Proper HA, van der Torre L. Changing Agreements: Intention Reconsideration based on Assumptions and Reasons. In: Ossowski S, Toni F, Vouros GA, editors. Proceedings of AT 2012. vol. 918 of CEUR Workshop Proceedings. CEUR-WS.org; 2012. p. 262-3.

[33] Zin MM, Satoh K, Borges G. Leveraging LLM for Identification and Extraction of Normative Statements. In: Proceedings of JURIX 2024. IOS Press; 2024. p. 215-25.

[34] Billi M, Pisano G, Sanchi M. Fighting the knowledge representation bottleneck with Large Language Models. In: Proceedings of JURIX 2024. IOS Press; 2024. p. 14-24.

[35] Pawar SS, Apte M, Palshikar GK, Ali B, Ramrakhiyani N. DRAssist: Dispute Resolution Assistance using Large Language Models. In: Proceedings ICAIL 2025. forthcoming. ACM; 2025. .

[36] Cabalar P, Ciabattoni A, van der Torre L. Deontic equilibrium logic with explicit negation. In: Proceedings of JELIA 2023; 2023. p. 498-514. Springer.

[37] Dong H, Doder D, Li X, Markovich R, van der Torre L, van Zee M. Rights and Practical Reasoning in Deontic Logic. In: Proceedings of DEON 2023. College Publications; 2023. p. 1-19.

[38] Yu L, Liga D, Markovich R. Addressing the Right to Explanation and the Right to Challenge through Hybrid-AI: Symbolic Constraints over Large Language Models via Prompt Engineering. In: Proceedings ICAIL 2025. forthcoming. ACM; 2025.