Heterogeneous Knowledge in Distributed Environmental Governance: Emergence, Integration, and the Preservation of Disagreement

Wanhong Huang
huangwanhong@serendip.ngo

Abstract

Governance of an interdependent environmental field requires knowledge that no party holds. Knowledge of such a field stands distributed among hydrological services, ecological programmes, health authorities, industrial operators, and populations resident in the systems concerned, and their holdings differ in variable, in resolution, in evidential standard, and in the kind of claim they issue. The prevailing response is integration understood as reconciliation: harmonise the categories, resolve the disagreements, and produce a single account on which decisions may rest. This paper argues that reconciliation so understood is unavailable and that its pursuit is harmful, and proposes an alternative conception of what integration can be. The method is conceptual, developed through an examination of what aggregation procedures preserve and what they discard. Three insights are advanced. Knowledge of a distributed field is not the union of what its parties hold, since the relations among their holdings, including their disagreements, carry information that no party's holding contains. Aggregation over claims that stand in logical relations to one another cannot in general be conducted claim by claim without producing conclusions no party accepts, which is a known impossibility whose bearing on environmental assessment has not been drawn. And the claims held by parties in such a field are of kinds that cannot be averaged, since a probability, a rate ratio, a monetary estimate, a qualitative account, and a legal duty are not five measurements of one quantity. From these the paper develops a conception of integration as temporary, question-relative assembly that preserves disagreement, uncertainty, and acknowledged ignorance as first-class content, together with the provenance required to prevent a single source counted twice from appearing as two confirmations. The paper surveys social epistemology, judgment aggregation, expert elicitation, post-normal science, co-production, epistemic injustice, and the traditional-knowledge literatures, states what each owns, and identifies the residue. It reports no empirical study and establishes no formal result.

Keywords: distributed knowledge; environmental assessment; judgment aggregation; disagreement; epistemic infrastructure

Notices

On the references. All 43 works cited in this paper have been checked against their sources at claim level: each source was consulted and confirmed to state what the citing sentence attributes to it. No entry in this paper's bibliography is unverified.
Status. This is a working draft circulated for discussion. It is a preliminary discussion paper and is not a finished statement of its author's position. Sections and statement numbers are subject to change.

Licence. This work is made available under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC 4.0).

Statement on the use of language models. Drafting, literature search and argumentative criticism for this paper were conducted in dialogue with large language models, specifically Claude (Anthropic) and ChatGPT (OpenAI). The claims, the structure, the selection of material and the position taken are the author's. References have been checked in stages, to two standards, and the process is incomplete. Part of the corpus has been verified at identity level, meaning author, title and publication coordinates were confirmed against the publisher of record. A smaller part has additionally been verified at claim level, meaning the cited source was consulted and confirmed to state what the citing sentence attributes to it. The remainder is unverified at either standard. The standard reached is marked on each bibliography entry: no mark for claim level, ‡ for identity level only, † for not yet verified. Corrections made so far are recorded in the project's citation-verification file; where a source proved not to bear the weight the text placed on it, the text was revised rather than the citation removed.

Companion papers. This paper is Paper 4 of a series on the governance of environmental change, and treats the assembly of distributed holdings with disagreement preserved as content. The other papers of the series treat, in order: the object; the field and its three structures; coordination through change; the event layer; criticality detection; observational capacity; representational requirements; routing and inquiry; the regime near a threshold; forecasting from the stream; the knowledge commons; attribution and ownership; responsibility at transitions; and supervisory duties. Paper 12 is reserved and is not yet drafted. A related paper by the author, developed independently of this series, treats the temporal correspondence limitation of evidence-based policy making, being the divergence between the system state that evidence represents and the system state at the moment of decision. Its results on the partition of a state space into estimable and unobservable components, on the blindness of a recency-weighted evidence base to the approach to criticality, and on the allocation of a justificatory burden where temporal correspondence fails, were reached by a different route and bear on several papers of this series.

Suggested citation. Huang, W. Heterogeneous Knowledge in Distributed Environmental Governance: Emergence, Integration, and the Preservation of Disagreement. Working draft.

Discussion Paper Note

This paper is a preliminary discussion paper intended to share an evolving idea and invite further dialogue, criticism, revision, and independent development.

The author does not claim exclusive epistemic ownership over the viewpoints, concepts, or lines of reasoning presented here, nor does the author claim priority as their first originator. Similar or related ideas may have appeared previously in other intellectual, cultural, or disciplinary traditions. Any legal rights retained in this work are intended to support attribution, responsible use, and protection against exploitative or harmful appropriation, and not to restrict independent inquiry, criticism, revision, or further development.

The arguments in this paper should therefore be understood as provisional and historically situated rather than definitive. Readers are encouraged to question, revise, extend, reinterpret, or independently develop the ideas presented here. Where appropriate, acknowledgment of this paper as one point of encounter in the development of related ideas is appreciated, but such acknowledgment should not be understood as granting the author epistemic ownership over the ideas themselves.

Responsible Use and Rights Reservation

The author encourages good-faith discussion, criticism, independent development, and responsible use of the knowledge presented in this work. The author does not claim exclusive epistemic ownership over the ideas or viewpoints discussed herein, nor claim priority as their first originator.

However, the author expressly reserves all rights and remedies available under applicable law with respect to uses of this work or related materials that involve unlawful conduct, harmful or abusive exploitation, improper commercial appropriation, infringement of applicable intellectual property or other legal rights, or conduct contrary to applicable national, regional, or international law.

This reservation is intended to preserve the ability to respond to misuse and harmful appropriation. It should not be interpreted as restricting legitimate academic inquiry, criticism, independent reasoning, or the further development of related ideas.

1 Introduction

A decision bearing on an interdependent environmental field requires knowledge that is distributed and remains so. A hydrological service holds records of flow; an ecological programme holds observations of populations and phenology; a health authority holds presentations and mortality; an operator holds performance data on the systems it runs; a resident population holds observation accumulated over generations and recorded in no register. None of these is a partial view of a single account that a sufficiently diligent compiler could assemble. They employ different variables, describe processes at different resolutions, apply different standards of evidence, and issue claims of different kinds.

The prevailing response treats this as a reconciliation problem. Categories are harmonised, ontologies are aligned, disagreements are adjudicated, and the product is a single assessed account on which decision may rest. Substantial institutional effort is directed at this, and the effort is not misdirected in every case: within a domain sharing a representation, harmonisation is achievable and useful. The paper's argument concerns integration across domains that do not share one, and its claim is that reconciliation in that setting is unavailable, that its pursuit consumes the effort available for coordination, and that the account it produces conceals what a decision-maker most needs to know.

Three considerations support this.

The first is compositional. The knowledge relevant to governing a distributed field is not the union of what its parties hold. That two independently constituted programmes disagree about a coupling is information about the field, and it is held by neither of them; that a party's account and a resident population's account diverge at a particular point locates a question that neither account raises. Relations among holdings are content, and a procedure that resolves them discards it.

The second is formal and is the paper's principal borrowing. Where parties hold views on claims that stand in logical relations to one another, aggregating their views claim by claim can yield a set of conclusions that no party holds and that may be inconsistent. This is established in the literature on judgment aggregation, and it is a general result about procedures and not a defect of any particular one. Environmental assessment routinely aggregates in exactly the vulnerable form, since its conclusions are conjunctions of premises on which parties are separately polled.

The third concerns the kinds of claim involved. A probability of exceedance, a rate ratio, a monetary damage estimate, an ethnographic account, and a statutory duty are not five measurements of a common quantity, and no weighting combines them. A procedure that produces a single figure from such inputs has not integrated them; it has selected one kind and expressed the others in its terms, and the selection is ordinarily made by whichever party's representation the procedure was built around.

Claim 1.1. *Knowledge adequate to govern a distributed environmental field is not held by any party and is not the union of what the parties hold. It is constituted by the relations among their holdings, including their disagreements, their uncertainties, and their acknowledged ignorance. An arrangement that integrates by resolving disagreement destroys part of what it was constituted to produce,*

Draft

and integration should therefore be understood as temporary, question-relative assembly that preserves disagreement as content.

Claim 1.1 is not a counsel of paralysis, and §4 treats the objection at length. An assembly that preserves disagreement still terminates, and it terminates in one of two ways: in an action taken on grounds that are stated, or in the finding that the uncertainty is too great to warrant the action contemplated. The second is a legitimate outcome and is ordinarily unavailable to an arrangement committed to producing a single account, since such an arrangement has no form in which to express it.

The paper proceeds as follows. §2 sets out what parties hold, why the holdings do not combine, and what assembly consists in. §3 surveys the literatures that own adjacent claims, which are numerous, and states the residue. §4 treats the objections, of which the decision objection is the most serious. §5 lists the open questions and §6 the declinations.

2 Holdings, Their Incombinability, and Assembly

2.1 Positions Carried from the Companion Papers

Three results from the earlier papers of this series are used here, and each is restated so that the present paper stands alone.

The first is the *object of governance*. The opening paper of the series argues that in a field whose dynamics exceed every party's representation, and in which every governing party is itself a member, governance cannot take outcomes or target states as its object, since the first presupposes a causal model no party holds and the second a shared state space heterogeneous parties do not have. The object that remains governable is the set of relational conditions under which trajectories are generated, comprising couplings, buffers, permissions, the distribution of access and information, and institutional rules. The bearing on the present paper is direct: the distribution of information is itself among the governed conditions, so the arrangement by which parties come to know a field is not preliminary to its governance but part of it.

The second is the *distinction among three structures*. The second paper distinguishes the physical structure of dependence in a field, which obtains independently of anyone's knowledge; the epistemic structure, recording which parties hold what knowledge of which dependencies, at what resolution and in which representational language; and the communicative structure, recording which parties in fact report what to which others. Their divergences are argued there to be governance-relevant quantities, and a dependence that obtains while no party holds knowledge of it is distinguished from one that is known and unreported. *The epistemic structure is the object of the present paper.* The second paper establishes that it diverges from the other two and treats it as a term in a comparison; here it is examined in itself.

The third is a *negative result and its consequence*. The second paper argues that assembling the observations of many parties into a single estimate of a common state presupposes a privileged representation of the system, which physics supplies to meteorology through its conservation laws and which coupled ecological, infrastructural, and social fields do not possess.

Draft

Integration of that kind is therefore unavailable across such a field's domains. The third paper supplies one positive counterpart, coordination through reported change, on the ground that agreement about the direction and unusualness of a change is far cheaper in representational commitment than agreement about the identity or value of the quantity that changed. The present paper supplies the other. If the parties' holdings cannot be reduced to a common representation, the question is what an arrangement should do with them unreduced, and that is the question taken up here.

2.2 Composition of a Party's Holding

A party's epistemic holding is more than a set of propositions, and the elements beyond the propositions are what determine whether it can be combined with another's.

Definition 2.1. *The holding of a party i is the tuple*

$$K_i = \big(C_i,\ M_i,\ E_i,\ U_i,\ S_i,\ P_i \big),$$

comprising the claims it advances, the method by which each was produced, the evidence on which each rests, the uncertainty attaching to each, the scope over which each is asserted, and the provenance recording where the underlying observation originated and through which hands it has passed. Two holdings are combinable with respect to a claim *where their claims concern a common object, their methods are commensurable, and their uncertainties are expressed on a common footing. Combinability is a relation between holdings and not a property of either.*

Four remarks on Definition 2.1 bear on what follows.

Scope is doing more work than it appears to. A claim asserted for a catchment under conditions observed in a particular decade is not the same claim as one asserted generally, and two parties may agree on the words while asserting them over different scopes. Assessment procedures that record conclusions and discard scopes produce agreement that dissolves on application.

Uncertainty is heterogeneous in kind. A confidence interval, a stated range of expert judgement, a qualitative expression of confidence, and an acknowledgement that a mechanism is not understood are four different objects, and only the first two are candidates for arithmetic. The last is the most consequential in an interdependent field and the least representable in the formats assessment procedures ordinarily employ.

Provenance is required for a reason beyond attribution. Where two parties report a finding that rests on the same underlying observation, their agreement is not two confirmations. Without lineage, an arrangement that weights claims by the number of parties advancing them systematically overweights well-circulated findings, and the effect is stronger where circulation is easy, which is the current condition.

Acknowledged ignorance belongs in the holding. A party's statement that it holds no account of a coupling is information, and it is information of a kind that assessment formats seldom carry, since a procedure organised around claims has no field for the absence of one.

Draft

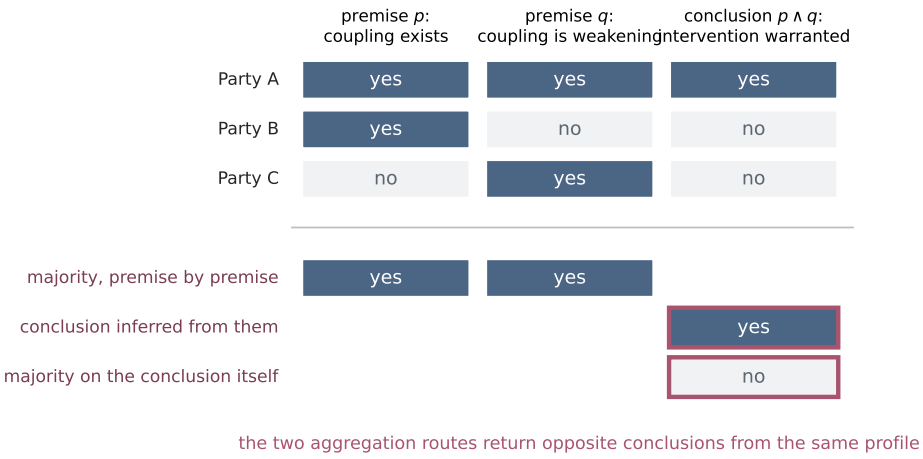

Figure 1: Aggregation over a conjunctive conclusion. Each party reasons consistently: the conclusion is asserted where both premises are asserted. Each premise is asserted by a majority. The conclusion is asserted by one party of three. A procedure polling premises returns the conclusion; a procedure polling the conclusion rejects it. The example is the discursive dilemma of the judgment-aggregation literature, presented here in the form environmental assessment routinely takes.

2.3 Obstacles to Combination

Three obstacles are distinguished, and they are of different orders.

Ontological. Parties describe the field in categories constituted by their own practices, and the categories are not always reconcilable. Soil moisture in an agronomic sense and in a hydrological sense are related and not identical, and the relation between them is not a conversion factor. Where the categories can be aligned, alignment is work and its product should carry an explicit statement of what the alignment assumed.

Evidential. Parties apply different standards for what counts as established. A discipline requiring controlled comparison and a discipline requiring sustained observation of a system that cannot be replicated will reach different verdicts on the same body of material, and neither is failing by its own lights. An arrangement that adopts one standard for combination has adjudicated a methodological dispute by procedural means.

Logical. Where claims stand in inferential relations, aggregation claim by claim is not in general coherent, and this obstacle is the sharpest because it is formal.

Figure 1 gives the structure, and the conditions under which the difficulty is general should be stated, since the paper's use of the result depends on them.

Let a set of parties each hold a judgement, accepting or rejecting, on every proposition in an agenda that is closed under negation and contains propositions standing in logical relations. Each party's judgement set is required to be consistent and complete over the agenda. An aggregation rule maps profiles of individual judgement sets to a collective judgement set. The impossibility result establishes that no rule satisfying the following three conditions delivers a consistent and complete collective set for every admissible profile, where the agenda is minimally connected in the sense the literature makes precise [26, 9, 27]:

- *Universal domain*: the rule accepts every profile of consistent and complete individual judgement sets.

Draft

- *Anonymity*: the rule treats the parties symmetrically, so that permuting the profile leaves the collective set unchanged.
- *Systematicity*: the collective judgement on each proposition depends only on the parties' judgements on that proposition, and depends on them in the same way for every proposition.

Three features of this statement bear on the use made of it here. The result concerns *procedures* and not the competence of parties; the parties in Figure 1 each reason without error. The escape routes are known and each abandons one condition: restricting the domain, abandoning anonymity by weighting parties, or abandoning systematicity by treating some propositions as premises and deriving the rest. And the third escape is the one assessment procedures ordinarily take without declaring it, since polling on components and composing is precisely the abandonment of systematicity in favour of a premise-based rule. Three parties are asked whether a coupling exists, whether it is weakening, and whether intervention is therefore warranted. Each party's answers are individually consistent. Both premises command majorities and the conclusion does not. A procedure that aggregates premise by premise and infers reaches a conclusion that two of three parties reject; a procedure that aggregates the conclusion directly reaches the opposite result. The disagreement between the two procedures is not resolvable by better data or more careful deliberation, and the choice between them determines the outcome.

Claim 2.2. *Environmental assessment aggregates in the form the impossibility results address. Its conclusions are ordinarily conjunctions, that a hazard exists, that exposure is material, and that a measure is effective, and its procedures ordinarily poll parties on components and compose the results. Where the parties are heterogeneous, the composed conclusion may be one that no party advances, and the aggregation procedure is therefore doing substantive work that is not visible in its output.*

Claim 2.2 does not assert that assessment is worthless, and §4 states the qualification. It asserts that the aggregation step carries a commitment, and that an arrangement should make the commitment visible in place of presenting its output as the parties' collective view.

2.4 Incommensurable Kinds of Claim

A further obstacle is distinct from the three above and is easily mistaken for a difficulty of measurement. The claims a governance question draws on are of different logical kinds. A probability of exceedance is a statement about a distribution. A rate ratio is a statement about a comparison between populations. A damage estimate is a statement about a valuation, itself resting on contested premises. An ethnographic account is a statement about meaning and practice. A statutory duty is not a description at all but a norm, and its force does not vary with the evidence for it.

No weighting combines these, and the appearance that one does is produced by expressing them all in the terms of one. Monetisation is the usual device and its effect is to convert the second and third into the terms of the fourth kind of object while leaving the fifth unrepresented, since a duty priced is a duty denied.

Draft

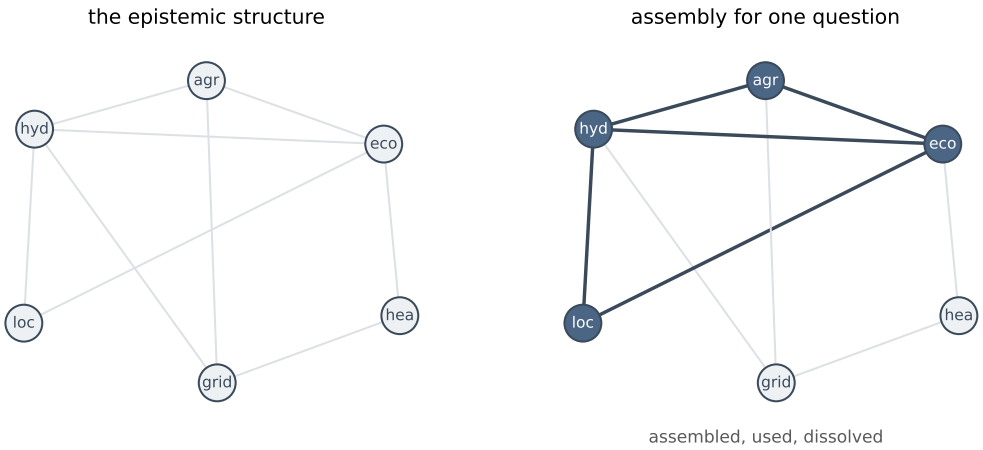

Figure 2: Question-relative assembly. The epistemic structure of a field relates many parties; a given question activates a subset of them and of the relations among them. The assembled structure is used and dissolved, and a different question activates a different subset. No party holds the whole, and no standing body maintains it.

The consequence for integration is that the assembled product must be able to carry claims of several kinds side by side, related to one another by stated relations of relevance and support, and not reduced to a common measure. The relations carry the integration; the reduction would destroy it.

2.5 Assembly, the Alternative to Reconciliation

The alternative to reconciliation is assembly, and the term is used in a definite sense.

Definition 2.3. *An* assembly *for a question q is a temporary structure $\mathcal{A}(q)$ drawing on the holdings of the parties whose knowledge bears on q, and comprising: the claims advanced, marked with their kind, scope, method, and provenance; the relations among them, of support, of tension, and of contradiction; the points at which the parties disagree, with the ground of disagreement stated where it is known; the uncertainties, in the forms in which they are held; and the acknowledged gaps, being the couplings and quantities for which no party advances a claim. An assembly is constructed for a question, is used in addressing it, and is dissolved; it is not a repository and is not maintained.*

Three properties of Definition 2.3 are load-bearing.

Assembly is *question-relative*. The parties whose knowledge bears on whether a reservoir operating rule should change are not those whose knowledge bears on whether a species is declining, and an arrangement that maintains one standing integrated account must either include everything, which is unattainable, or fix a scope in advance, which prejudges which questions can be asked.

Assembly is *temporary*. A standing integrated account acquires authority with age and becomes the object that subsequent work is required to reconcile with, which converts an assessment into an institution. A dissolved assembly leaves the holdings as it found them, and the next question is addressed from the parties' current holdings and not from the residue of the last question.

Draft

Assembly *terminates in one of two ways*, and the second is what distinguishes it. It terminates in action taken on grounds that are stated, including the disagreements the grounds did not resolve, or in the finding that the assembled material does not warrant the action contemplated. The second outcome is a finding and not a failure, and an arrangement that cannot express it will convert every assembly into a warrant.

2.6 Formal Apparatus for an Assembly

Definition 2.3 states what an assembly contains and leaves its structure unspecified, which is a defect if the assembly is to support operations. Two established formalisms supply what is needed, and the paper adopts them without constructing anything of its own.

The relational structure. An assembly's claims and the relations among them form an argumentation framework in the sense of Dung [11]: a set of arguments together with a binary attack relation. The semantics of such frameworks determine which sets of arguments are collectively admissible, and several notions of admissibility are defined, of which the relevant feature here is that a framework with mutually attacking arguments ordinarily admits *several* maximal admissible sets and not one. Applied to an assembly, the several extensions are the several defensible positions the material supports, and the existence of more than one is the formal counterpart of preserved disagreement. Where extensions are supported by claims of different kinds, the structure additionally records what each position rests on.

The quantitative structure. Where parties advance probabilistic claims that do not agree, the pooled point estimate of §3 is one option and the set of the parties' distributions is another. The theory of imprecise probability [41] treats a convex set of distributions, a credal set, as the primitive object, from which lower and upper expectations follow; the theory of belief functions [36] distinguishes the support committed to a proposition from that committed to its negation, leaving a remainder that represents commitment to neither. Both formalisms carry a distinction the pooled estimate cannot: between a quantity known to be uncertain within stated bounds and a quantity for which the parties' commitments leave a range unassigned. This is the distinction §2 requires between uncertainty and acknowledged ignorance.

Claim 2.4. *An assembly admits a formal representation in which the claims and their relations of support and attack constitute an argumentation framework, and the quantitative claims constitute a credal set or a belief function in place of a point estimate. Under this representation, preserved disagreement is not an absence of structure: multiple admissible extensions are a determinate output, the width of the credal set is a determinate quantity, and the mass committed to neither a proposition nor its negation measures acknowledged ignorance. An assembly therefore supports operations, and the operations return sets where a consensus procedure returns points.*

Claim 2.4 is a claim about the availability of a representation and not a formal result, and three limits on it are recorded. Argumentation semantics are defined over attack relations that must be supplied, and the identification of attacks among claims stated in different representations is exactly the difficulty §2 describes; the formalism organises the difficulty and does not dissolve it. Credal sets and belief functions require that the quantitative claims concern a common proposition, so they apply within the combinable subset of an assembly and

Draft

	Reconciliation to a consensus	Question-relative assembly
Output	A single account, with an uncertainty	A structure of claims, relations, and gaps
Kinds of claim admitted	Those expressible in the adopted representation	Several, carried with their kinds
Treatment of disagreement	Resolved, or recorded as a residual	Retained with its ground, as content
Treatment of ignorance	Ordinarily unexpressed	An unassigned commitment, reported
Where the choice is made	Inside the aggregation rule	By the accountable party, visibly
Standing over time	Accrues authority; later work reconciles to it	Dissolved after use
Characteristic failure	A conclusion no party advances, undiscoverable from the output	Exploitation by a party manufacturing doubt

Table 1: Two conceptions of integration compared. The final row states each conception's characteristic failure, and §4 argues that the second fails visibly where the first fails invisibly.

not across it. And no operation on the combined structure, argumentative and quantitative together, is supplied by either literature, which §5 records as an open question.

2.7 The Standing of the Union Claim

Claim 1.1 asserts that collective knowledge is not the union of the parties' holdings, and the assertion should be made precise, since a union of sets of claims is a well-defined object and the claim is not that it fails to exist.

Let $K_1, \ldots, K_n$ be the holdings and let $\bigcup_i K_i$ be the set of claims any party advances. Three kinds of content are governance-relevant and absent from that set.

Relational content. That K_i and K_j disagree at a stated point is a fact about the pair, and it is a member of neither. Where the disagreement is grounded, the ground is likewise a fact about the pair. The union of the holdings contains the claims and not the relations among them, and §2 argues the relations are what a decision-maker requires.

Comparative content. That a claim advanced by one party is supported by an independently constituted party is different from its being repeated by a party drawing on the same observation, and neither fact is a member of the union, since both concern provenance relations across holdings.

Negative content. That no party advances a claim on a coupling is a fact about the set of holdings taken together, and a union contains no record of what its members omit.

The precise statement is therefore that governance-relevant knowledge is a structure over the holdings and not a set of claims drawn from them, and that the union operation discards the structure. This is weaker than the rhetorical form of Claim 1.1 and is what the paper argues.

2.8 Worked Case: A Reservoir Rule under Contested Evidence

The following traces an assembly through a case, and is offered to show what the structure of §2 contains and what it returns.

Draft

A reservoir authority proposes to alter its operating rule, retaining more water through the early season against a drier summer. Five parties hold relevant knowledge. The authority holds operational records and a calibrated hydrological model. A downstream ecological programme holds twenty years of observation on a fish population whose spawning depends on a spring flow pulse. An agricultural body holds abstraction records and yield data. A health authority holds evidence on the exposure of a settlement to summer heat, which the reservoir's cooling function affects. A resident community holds observation of the river's behaviour across a longer period than any register covers.

The claims are of several kinds. The authority advances a probability of shortfall under each rule. The ecological programme advances a claim about a causal dependence, with a confidence expressed qualitatively. The agricultural body advances a valuation. The health authority advances a rate ratio conditional on exposure. The resident community advances an account of the river's past behaviour that is detailed, not quantified in the terms the others use, and covering years for which no gauged record exists.

The assembly records these with their kinds, scopes, methods, and provenance, and then records the relations among them. Three findings result.

The first is a located disagreement. The authority's model and the resident account differ about the frequency of a low-flow condition, and the ground of the difference is identifiable: the model is calibrated on the gauged record, which begins after the period the resident account describes. This is not a disagreement about facts within a shared frame but a disagreement about the period over which the frame is valid, and locating it tells the decision-maker what the model's confidence interval does not cover.

The second is an unassigned commitment. No party advances any claim about the interaction between the retained volume and the settlement's summer heat exposure, though two parties hold claims about each side of it separately. The gap is a determinate output, and it is the kind of output no consensus figure carries.

The third is a conjunctive vulnerability of the sort Claim 2.2 describes. The proposition that the new rule is warranted decomposes into a claim about shortfall risk, a claim about ecological cost, and a claim about their relative weight. Polling the parties on the components and composing returns a warrant; asking each party whether the rule is warranted, on its own view of all three, returns a majority against. The assembly records both, and the discrepancy is itself reported.

The assembly returns a structure with several admissible positions, a bounded range in place of a point estimate of shortfall risk, one located disagreement with its ground, one identified gap, and a note that the two aggregation routes diverge. It returns no recommendation. The decision remains with the authority, and Claim 4.1 states what has changed: the authority now chooses visibly among positions the material supports, in place of receiving a figure in which the choice has already been made.

Draft

2.9 The Informational Content of Disagreement

The proposal that disagreement be preserved requires a positive account of what preservation yields, since otherwise it is only a refusal.

Disagreement locates. Two parties agreeing about a system's behaviour and disagreeing about a particular coupling have identified the coupling as the place where the field is not understood, and the identification is available to neither alone.

Disagreement discriminates among kinds of doubt. A disagreement grounded in different data is answerable by observation; one grounded in different methods is answerable, if at all, by methodological argument; one grounded in different values is not an epistemic matter and should not be presented as one. An arrangement that records the ground of a disagreement supplies a decision-maker with information about what further work would resolve it, and a resolved account supplies none.

Disagreement bounds. A range of claims advanced by independently constituted parties is evidence about how wide the plausible range is, and the width is a quantity a decision-maker needs. A procedure that reports a consensus figure with an uncertainty derived from within one method understates it, and understates it systematically, since the neglected variance lies between methods and not within them.

Persistent disagreement is itself a finding. Where parties with adequate evidence continue to disagree after their grounds have been examined, the arrangement has established that the question is not settled by the available evidence, and that is a determinate result which a forced consensus conceals.

3 Prior Treatments of Distributed and Aggregated Knowledge

Nine literatures bear on the position. The concessions are extensive, and the residue is correspondingly narrow.

3.1 Aggregation of Judgment and Its Impossibility

The formal result on which §2 rests belongs to the judgment-aggregation literature. List and Pettit [26] established that no aggregation procedure satisfying a small set of natural conditions, universal domain, anonymity, and systematicity, delivers consistent collective judgments over a logically connected agenda, and subsequent work has mapped the conditions under which escape is possible [27, 9]. The phenomenon was identified in the legal context as the doctrinal paradox, where a court reaching a verdict by aggregating over premises may reach the opposite of the verdict its members reach individually [24]. The relation to Arrow's theorem on preference aggregation [2] is now well understood, and both belong to a family of impossibility results about the aggregation of structured attitudes.

The concession is complete: the result is theirs, its conditions are theirs, and this paper contributes nothing to it. The paper adds an application. The literature has been developed with courts, committees, and expert panels as its illustrations, and its bearing on environmental assessment, whose conclusions are conjunctive and whose procedures poll on components,

Draft

has not to the author's knowledge been drawn out. Where such a drawing-out exists in a literature not surveyed here, priority belongs to it.

3.2 Pooling of Opinions and Probabilities

A parallel literature addresses the combination of probability distributions held by several parties. Linear and logarithmic pooling rules and their axiomatic characterisations are established [15, 37], and Bayesian approaches treat other parties' judgements as evidence to be conditioned on [31]. Structured expert judgement supplies operational procedures for weighting experts by measured performance on calibration questions [8].

These methods are effective under a condition they require: the parties' judgements must concern a common quantity, expressed on a common scale. Where the condition holds, the present paper's argument does not apply and pooling should be used. §2 concerns the case where it does not hold, and the concession is therefore about scope: pooling owns the combinable case, and the paper addresses the remainder.

Two findings from this literature transfer to the remainder and are used. Dependence among sources degrades pooled estimates in ways that the pooled uncertainty does not display, which is the formal counterpart of the provenance requirement in Definition 2.1. And performance-based weighting requires calibration questions with known answers, which are ordinarily available for short-horizon quantities and unavailable for exactly the long-horizon structural claims that environmental governance turns on.

3.3 Social Epistemology and the Epistemic Value of Diversity

Social epistemology treats knowledge as produced by communities and examines the arrangements that make communities reliable [16, 28]. Longino's account makes the availability of criticism from differently situated inquirers a condition of objectivity, which is an argument for preserving disagreement in place of resolving it and is a direct ancestor of the position here. Formal work on the epistemic properties of groups establishes that diversity of problem representations can outperform individual ability under stated conditions [19], and the Condorcet tradition establishes conditions under which majorities track truth, together with the dependence of that result on the independence of the judgements [17].

These literatures own the claim that a plurality of situated inquirers is epistemically valuable and not merely politically desirable, and the concession is unreserved. The residue is narrow: they establish that diversity is valuable and address the design of communities of inquiry, and the present paper asks what a governing arrangement should do with a body of unreconciled holdings at the point of decision, which is a question about procedure and not about the epistemic virtues of communities.

3.4 Post-Normal Science and Extended Peer Review

Funtowicz and Ravetz [14] characterised a class of problems in which facts are uncertain, values are in dispute, stakes are high, and decisions are urgent, and argued that quality assurance in such problems requires an extended peer community including parties outside the

Draft

credentialed disciplines. The associated work on uncertainty in policy-relevant science distinguishes technical uncertainty from methodological and epistemological uncertainty and from ignorance proper [43, 39], and the last of these is the category that Definition 2.1 carries as acknowledged ignorance.

The concession is substantial: this tradition owns the analysis of uncertainty's kinds and the argument for extended participation in quality assurance. Two differences from the present position should be recorded. Post-normal science addresses the quality of a scientific product entering a policy process, and the present paper addresses the structure of an assembly used in reaching a decision, which is a later moment. And the extended peer community is a body constituted for review, where an assembly under Definition 2.3 is constituted for a question and dissolved after it.

3.5 Co-Production and the Politics of Assessment

Jasanoff's account of co-production [21] holds that the natural and social orders are produced together, so that a scientific assessment and the political order it serves shape one another. The associated work on regulatory science and on civic epistemologies [22] establishes that what counts as adequate evidence differs across polities in ways that are not reducible to differences in scientific capacity. Studies of the assessment process in climate science [12, 3] document the procedures through which a heterogeneous literature is rendered into consensus statements, including the treatment of dissent.

This literature owns the finding that assessment procedures are consequential and contested, and that the production of a single account is a political act as much as an epistemic one. The present paper takes that finding as a premise and asks what follows for the design of an alternative, which the literature does not undertake, its interest being analytic and critical.

One specific finding is used. The documented practice of consensus assessment includes the systematic narrowing of reported uncertainty as material passes upward through summarisation, which is the effect Claim 2.2 predicts and which §2 attributes to the aggregation step.

3.6 Epistemic Injustice and the Standing of Testimony

Fricker [13] distinguished testimonial injustice, in which a speaker is accorded less credibility on account of identity-based prejudice, from hermeneutical injustice, in which a party lacks the interpretive resources to render its experience intelligible. Dotson [10] extended the analysis to forms of silencing that operate through the audience's unwillingness or inability to receive. The bearing on the present setting is direct, since the parties in an environmental field include populations whose accounts are systematically discounted.

The concession is complete and the paper takes the analysis as it stands. One qualification is recorded because it affects what remedy is available. Where a party's account is discounted because it is not expressed in the form an assessment procedure receives, the deficit is procedural and the remedy is a procedure that receives other forms, which Definition 2.3 is designed to provide. Where a party's account is discounted because it rests on a shorter or

Draft

thinner record, the deficit tracks a real difference in evidential support, and correcting prejudice leaves it in place; the remedy for that case lies in the distribution of observational capacity and not in the reception of testimony.

3.7 Traditional and Local Ecological Knowledge

A substantial literature establishes that populations resident in ecological systems hold knowledge of them that is detailed, dynamic, and not reducible to the categories of scientific ecology [6, 1]. Work on the integration of such knowledge with scientific assessment documents both its value and the characteristic failure of integration attempts, in which local knowledge is admitted as data to be validated by scientific method and thereby stripped of the framework that made it intelligible [33].

This literature owns the substantive claim and, importantly for the present paper, owns the diagnosis of the failure mode. Nadasdy's analysis of integration as a process that distils local knowledge into forms compatible with the receiving institution is precisely the failure that Definition 2.3 attempts to avoid by preserving claims in the kind and scope in which they are held. The paper's contribution here, if any, is to state a general structure of which that failure is an instance, and the priority of the diagnosis is conceded.

3.8 Distributed Cognition and Multi-Agent Architectures

Hutchins [20] established that cognitive processes may be properly attributed to systems comprising several people and their instruments, and that the analysis of such systems requires attention to the propagation of representational states across media. The engineering counterpart, in blackboard architectures and multi-agent systems, addresses the coordination of specialist components that hold partial and heterogeneous representations of a problem [34, 42], and the literature on federated and distributed inference addresses the combination of locally held models without centralising data [30].

The concession is that the general possibility of a cognitive process distributed across heterogeneous components is established. The differences bearing on the present setting are that the components in these treatments are designed to interoperate, whereas the parties in a governance field were constituted independently and for other purposes; and that a designer stands outside the multi-agent system and specifies its protocol, whereas a governance arrangement is constituted by the parties it coordinates.

3.9 Formal Treatments of Inconsistent and Imprecise Bodies of Claim

Two formalisms adopted in §2 have literatures of their own, and the concessions are unreserved.

Dung's theory of abstract argumentation [11] takes a set of arguments and an attack relation as primitive and defines several notions of collective acceptability, establishing that a framework may admit multiple extensions. Subsequent work has developed structured argumentation, in which arguments are constructed from a knowledge base and attacks derived

Draft

from its inconsistencies [35, 7], and the field has been applied to deliberation and to decision support. The relevance to the present paper is that this literature supplies the operations that Definition 2.3 requires and does so for exactly the case of a body of claims that is not consistent.

Imprecise probability [41] and the theory of belief functions [36] supply representations of partial commitment. Robust Bayesian analysis and the associated sensitivity literature address the practical treatment of sets of priors [5]. The distinction these formalisms carry between uncertainty and unassigned commitment is the one the present paper requires for acknowledged ignorance, and its priority is theirs.

The paper contributes nothing to either formalism and adopts both. The proposal here is their conjoint use in a governance setting, with argumentation carrying the relations among heterogeneous claims and imprecise representations carrying the quantitative ones, and §5 records that the combination is not supplied by either literature.

3.10 Ensembles and Model Plurality in Environmental Science

Climate and environmental science have confronted the problem of many models directly. Multi-model ensembles are assembled from models developed independently, and the resulting spread is used as an indication of uncertainty [38, 23]. The literature is candid about the difficulties, and three of them bear on the present paper.

The first is that such collections are ensembles of opportunity, comprising the models that exist and not a sample from a defined population, so the spread has no straightforward statistical interpretation. The second is that models are not independent: they share components, parameterisations, and intellectual ancestry, so agreement among them overstates confirmation, which is the provenance problem of Definition 2.1 appearing in the modelling domain. The third is the standing dispute between model democracy, in which each model counts equally, and performance weighting, in which models are weighted by skill against observation, with the objection that skill on observable quantities may not indicate skill on the projected ones.

This literature owns the analysis of model plurality and its pathologies, and priority is conceded. The bearing on the paper is twofold. It supplies independent evidence for the dependence problem the provenance requirement addresses. And its unresolved dispute about weighting is the same question that arises for parties in an assembly, arrived at in a domain where the entities are models and the performance record is comparatively rich; the difficulty is harder in the present setting, where the entities are institutions and no calibration record exists.

3.11 Calibrated Language and the Reporting of Agreement

Assessment practice has developed apparatus for reporting confidence that is closer to the present proposal than any other institutional instance. The guidance developed for the Intergovernmental Panel on Climate Change specifies a calibrated vocabulary in which a finding's confidence is expressed qualitatively as a function of two dimensions, the evidence, char-

Draft

acterised by its type, amount, quality and consistency, and the degree of agreement, while likelihood is expressed separately on a quantified probabilistic scale [29, 32]. The guidance is explicit that agreement is not equivalent to consistency, which is a distinction of the kind §2 argues an assembly must carry: consistency is a property of the evidence, and agreement is a property of those assessing it, so evidence may be consistent and its assessors divided, or inconsistent and its assessors united in reading it the same way.

The two-dimensional treatment is the important element. A finding supported by ample evidence with low agreement is distinguished from one supported by limited evidence with high agreement, and the two are not collapsed into a single confidence figure. This is disagreement carried into the output as content, and priority for the device belongs to that practice.

Two differences from the present proposal are recorded. The apparatus reports agreement as a summary attribute of a finding and does not carry the parties, their grounds, or the locations of their disagreement, so a reader learns that agreement is low and not who disagrees or why; and its own distinction between agreement and consistency, having been drawn, is not carried into the output, since a single confidence term is reported and the reader cannot recover which of the two dimensions produced it. And the apparatus operates on lines of evidence within a broadly shared scientific representation, and the parties of §2 include holders of claims of other kinds entirely. The proposal here may be read as an extension of the calibrated-language device to a setting with heterogeneous claim kinds and with the grounds of disagreement retained.

3.12 Decision under Deep Uncertainty

A body of work addresses decision where probabilities cannot be agreed, and it is the closest existing answer to the decision objection of §4. Robust decision making [25] searches for strategies that perform acceptably across a wide range of plausible futures in place of optimising against a single distribution, and the associated field of decision making under deep uncertainty develops adaptive pathways and related methods [40, 18]. Scenario analysis supplies plural futures where a forecast is unavailable, and info-gap theory [4] treats the robustness of a decision to error in its own model as the object of analysis.

The concession is substantial and the relation is complementary. That literature takes a plurality of futures as given and asks what decision procedure is appropriate; the present paper asks how a plurality of accounts held by institutionally distinct parties should be constituted and carried, which is prior to the decision procedure and is not treated there. The two fit: an assembly preserving disagreement is a natural input to a robustness analysis, and a robustness criterion is a natural answer to the objection that an assembly returns no single figure. §5 records the connection as a question, since the fit has not been worked out.

3.13 Residue of the Survey

Claim 3.1. *Consider parties independently constituted, holding claims of several logical kinds over incompatible categories, standing in no hierarchical relation, addressing a question whose conclusion*

Draft

is a conjunction of components on which they may be separately polled, in a setting where no calibration questions with known answers are available and no designer stands outside. The claim that integration for such parties should take the form of temporary, question-relative assembly preserving disagreement, uncertainty, kind, scope, and provenance as first-class content, and terminating either in action on stated grounds or in the finding that the material does not warrant the action, is not made by the treatments surveyed. Judgment aggregation supplies the impossibility and addresses procedures for reaching collective judgments; pooling supplies methods for the combinable case; social epistemology establishes the value of diversity and addresses communities of inquiry; post-normal science supplies the taxonomy of uncertainty and addresses quality assurance; co-production supplies the critique and declines the constructive task; epistemic injustice supplies the analysis of discounted testimony; the traditional-knowledge literature supplies the diagnosis of integration's characteristic failure in one domain; and distributed cognition supplies the possibility under conditions of designed interoperability that the present setting lacks.

4 Objections and the Limits of the Position

4.1 The Objection from Decision

The most serious objection is that decisions must be made, that a decision requires a determinate input, and that an arrangement returning a structured disagreement has evaded the task it was constituted for. A minister requires an answer; an assembly that reports three incompatible accounts and their grounds has supplied none.

The objection has force and is answered in three parts.

The first part concedes the premise and denies the inference. A decision requires a determination and does not require a determinate epistemic input, and the two are ordinarily conflated. A decision-maker choosing under acknowledged disagreement is choosing, and choosing in a way that can be stated and reviewed. The objection demands that the epistemic apparatus absorb the burden of choice by returning a single account, and the effect is that the choice is made inside the assessment procedure, by the aggregation rule, where it is neither visible nor attributable.

The second part concerns what the single account costs. Claim 2.2 establishes that composing a conjunctive conclusion from separately polled components can return a conclusion the parties reject. A decision-maker receiving such a conclusion is not receiving the parties' view and has no means of discovering this from the output. The consensus is more determinate and less informative, and the determinacy is manufactured by the procedure.

The third part states what an assembly delivers that is decision-relevant. It delivers the range of accounts held by parties competent to hold them, which bounds the plausible outcomes more honestly than a within-method uncertainty; the location of the disagreements, which identifies where the field is not understood; the ground of each disagreement, which indicates what further work would resolve it and whether the disagreement is epistemic at all; and the acknowledged gaps, which a consensus format cannot express. A decision-maker holding these is better placed than one holding a figure, and is worse placed only in the re-

Draft

spect that the choice is visibly theirs.

Claim 4.1. *An assembly preserving disagreement does not withhold a decision and relocates it. The choice among incompatible accounts is made by the party accountable for the decision, on grounds that can be stated and reviewed, in place of being made by an aggregation rule inside the assessment procedure, where it is neither visible nor attributable. This is the position's principal institutional consequence and its principal cost, since it removes from decision-makers a resource they value, which is the appearance that the choice was determined by the evidence.*

4.2 The Objection from Manipulation

An arrangement that preserves disagreement is exploitable by a party with an interest in delay. Manufacturing a disagreement is cheaper than resolving one, and the history of contested regulation supplies abundant instances of interests sustaining apparent scientific dispute to defer action.

The paper does not have a full answer and states what it has. Definition 2.3 requires that a disagreement carry the ground on which it rests, and a manufactured disagreement is characteristically weak in exactly that respect, since its ground is either absent or reduces on examination to a claim about what has not been proven. An arrangement recording grounds therefore distinguishes disagreements that locate a real question from those that do not, in a way that a bare tally of dissent does not. This is a partial answer, because assessing a ground requires competence that the assessing party may lack, and because a well-resourced interest can supply grounds of apparent quality.

Two further observations bear. The manipulation risk is not avoided by consensus formats, which are exploitable at the point where the consensus is drafted and where a single sustained objection can weaken a statement without appearing in the output. And the risk is asymmetric between the two failure modes: an arrangement that suppresses genuine disagreement fails invisibly, and one that admits manufactured disagreement fails visibly and is therefore correctable. The paper takes the second to be preferable and acknowledges that this is a judgement.

4.3 The Objection from Cost

Assembly is expensive. Recording kind, scope, method, uncertainty, provenance, and the grounds of disagreement, for each question, requires effort that a consensus figure does not, and the effort is required of parties that have other duties.

The paper concedes the cost and observes that the comparison is ordinarily made against an idealisation. The alternative is not a cheap consensus but a consensus purchased through prolonged negotiation over categories and wording, and the cost of that negotiation is borne before any question is addressed. Whether assembly is cheaper in aggregate is an empirical question the paper does not resolve, and the honest statement is that it redistributes cost from the harmonisation of categories to the documentation of particular questions.

Draft

4.4 Limits Internal to the Proposal

Four limits are internal and are stated as such.

The proposal supplies no rule for when an assembly is adequate. It states that an assembly terminates in action or in the finding that action is unwarranted, and it does not state how the second is reached, which leaves a determination of considerable consequence unspecified.

The proposal presupposes that parties can articulate the grounds of their disagreements. Some disagreements rest on judgement acquired through practice that its holder cannot render explicit, and an arrangement requiring articulated grounds will systematically underweight them, which reproduces at a second remove the exclusion that Definition 2.3 was designed to avoid.

The proposal treats the set of parties as given. Who is assembled is a determination prior to everything the assembly does, and the paper supplies no account of how it is made or of what makes an omission an error.

The proposal has no answer to a field in which the holdings are not merely different but adversarial, in which parties advance claims strategically and provenance is falsified. The machinery of Definition 2.1 presupposes good-faith reporting of method and evidence, and an arrangement operating without that presupposition requires verification the paper does not describe.

5 Questions Left Open for the Programme

5.1 Questions Concerning Structure

Open Question 5.1. ***Representation of an assembly.*** *Definition 2.3 specifies what an assembly contains and not how it is represented. An answer would supply a structure capable of carrying claims of several logical kinds, relations of support and tension among them, uncertainties in several forms, and acknowledged gaps, and would state what operations over that structure are well defined, since a structure supporting no operations is a filing convention.*

Open Question 5.2. ***Escape from the impossibility.*** *The judgment-aggregation literature identifies conditions under which consistent aggregation is available, including restrictions on the agenda and relaxations of systematicity. An answer would establish whether the agendas arising in environmental assessment satisfy any of the escape conditions, and if some do, would characterise them, since the result would then apply only to the remainder.*

Open Question 5.3. ***Relations among claims of different kinds.*** *The proposal requires that a probability, a valuation, an ethnographic account, and a duty stand side by side related by stated relations. An answer would state what those relations are and what inferences they license, since relations that license no inference leave the assembly inert.*

Open Question 5.4. ***Combination of the two formalisms.*** *Claim 2.4 adopts an argumentation framework for the relational structure and an imprecise representation for the quantitative one, and neither literature supplies operations over the combination. An answer would state how a credal set*

Draft

attached to a claim participates in an attack relation, and what an extension of the combined structure consists in.

Open Question 5.5. ***Derivation of attacks across representations.*** *Argumentation semantics presuppose a supplied attack relation, and identifying attacks among claims held in incompatible representations is the difficulty the paper attributes to integration generally. An answer would state what evidence establishes that two claims held in different representations are in conflict, as distinct from being about different things.*

Open Question 5.6. ***Interpretation of extension multiplicity.*** *A framework admitting several extensions is the formal counterpart of preserved disagreement, and the number of extensions is sensitive to the granularity at which claims are individuated. An answer would state whether the multiplicity is a stable feature of a body of claims or an artefact of its formalisation.*

5.2 Questions Concerning Practice

Open Question 5.7. ***Determination of sufficiency.*** *An assembly terminates in action or in the finding that the material does not warrant it. An answer would state what makes the second determination correct, and would confront the asymmetry that the finding is easy to reach for a party opposing the action and hard for one favouring it.*

Open Question 5.8. ***Elicitation of grounds.*** *Preserving a disagreement requires eliciting the ground on which it rests, and grounds are not always available to their holders. An answer would state what elicitation procedures recover grounds from practical judgement, and would say how an arrangement should treat a disagreement whose ground cannot be articulated, which is a case the proposal currently disadvantages.*

Open Question 5.9. ***Constitution of the party set.*** *An answer would state how the parties assembled for a question are determined, what makes an omission an error, and who is competent to find that one has occurred. This question is prior to the rest of the proposal and is presently unaddressed.*

Open Question 5.10. ***Verification under adversarial conditions.*** *An answer would state what verification of method, evidence, and provenance is feasible where reporting is strategic, and would characterise the classes of claim for which verification is available and those for which it is not.*

Open Question 5.11. ***Composition with robustness criteria.*** *Decision under deep uncertainty supplies criteria that operate over a plurality of futures, and an assembly supplies a plurality of accounts. An answer would state whether the two compose, and specifically whether a robustness criterion can operate over a set of admissible positions that differ in the kinds of claim they rest on.*

5.3 Questions Concerning Consequence

Open Question 5.12. ***Effect on decision quality.*** *The paper argues that a decision-maker holding a structured disagreement is better placed than one holding a consensus figure, and offers no evidence.*

Draft

An answer would state what evidence bears on the comparison and would confront the difficulty that the counterfactual is unobservable in any particular case.

Open Question 5.13. ***Accountability under acknowledged uncertainty.*** *Claim 4.1 relocates a choice to the accountable party. An answer would state what standard of review applies to a choice made under acknowledged disagreement, since a reviewer assessing such a choice against outcomes will penalise honest uncertainty, and a reviewer assessing it against grounds requires a standard the paper has not supplied.*

Open Question 5.14. ***Interaction with legal evidentiary standards.*** *Legal proceedings apply standards of proof to which an assembled structure is not obviously responsive. An answer would state how a structure preserving disagreement enters a proceeding requiring a finding, and whether the preservation survives that entry or is undone by it.*

Open Question 5.15. ***Scale of application.*** *The examples in this paper concern questions of moderate scope. An answer would state whether assembly is available for questions involving very many parties, and would confront the possibility that the proposal is workable at the scale of a basin and unworkable at the scale of a planetary assessment, which would confine it substantially.*

6 Declinations and Limits of the Position

The paper establishes no formal result. The impossibility on which §2 rests belongs to the judgment-aggregation literature, and the paper draws a consequence from it for a setting that literature has not addressed.

The paper claims no novelty in the analysis of uncertainty's kinds, which belongs to post-normal science, nor in the critique of consensus assessment, which belongs to science and technology studies, nor in the diagnosis of integration's failure for local knowledge, which belongs to the traditional-knowledge literature.

The paper opposes no method that operates within its conditions. Where parties hold judgements about a common quantity on a common scale, pooling and structured expert judgement are appropriate and the argument does not reach them.

The paper supplies no representation, no procedure, and no institution. It states what an assembly must contain and declines to specify the format in which it is carried, the procedure by which it is constructed, or the body that would convene it, and §5 records these as the programme's obligations and not as details.

References

[1] Agrawal, A. (1995). Dismantling the divide between indigenous and scientific knowledge. *Development and Change*, 26(3), 413–439.

[2] Arrow, K. J. (1951). *Social Choice and Individual Values*. Wiley.

Draft

[3] Beck, S. (2011). Moving beyond the linear model of expertise? IPCC and the test of adaptation. *Regional Environmental Change*, 11, 297–306.

[4] Ben-Haim, Y. (2006). *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edn. Academic Press.

[5] Berger, J. O. (1994). An overview of robust Bayesian analysis. *Test*, 3(1), 5–124.

[6] Berkes, F., Colding, J., and Folke, C. (2000). Rediscovery of traditional ecological knowledge as adaptive management. *Ecological Applications*, 10(5), 1251–1262.

[7] Besnard, P. and Hunter, A. (2008). *Elements of Argumentation*. MIT Press.

[8] Cooke, R. M. (1991). *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press.

[9] Dietrich, F. and List, C. (2007). Arrow's theorem in judgment aggregation. *Social Choice and Welfare*, 29(1), 19–33.

[10] Dotson, K. (2011). Tracking epistemic violence, tracking practices of silencing. *Hypatia*, 26(2), 236–257.

[11] Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2), 321–357.

[12] Edwards, P. N. (2010). *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming*. MIT Press.

[13] Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

[14] Funtowicz, S. O. and Ravetz, J. R. (1993). Science for the post-normal age. *Futures*, 25(7), 739–755.

[15] Genest, C. and Zidek, J. V. (1986). Combining probability distributions: a critique and an annotated bibliography. *Statistical Science*, 1(1), 114–135.

[16] Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford University Press.

[17] Grofman, B., Owen, G., and Feld, S. L. (1983). Thirteen theorems in search of the truth. *Theory and Decision*, 15(3), 261–278.

[18] Haasnoot, M., Kwakkel, J. H., Walker, W. E., and ter Maat, J. (2013). Dynamic adaptive policy pathways. *Global Environmental Change*, 23(2), 485–498.

[19] Hong, L. and Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46), 16385–16389.

[20] Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

[21] Jasanoff, S. (ed.) (2004). *States of Knowledge: The Co-Production of Science and Social Order*. Routledge.

[22] Jasanoff, S. (2005). *Designs on Nature: Science and Democracy in Europe and the United States*. Princeton University Press.

[23] Knutti, R., Furrer, R., Tebaldi, C., Cermak, J., and Meehl, G. A. (2010). Challenges in combining projections from multiple climate models. *Journal of Climate*, 23(10), 2739–2758.

Draft

[24] Kornhauser, L. A. and Sager, L. G. (1993). The one and the many: adjudication in collegial courts. *California Law Review*, 81(1), 1–59.

[25] Lempert, R. J., Popper, S. W., and Bankes, S. C. (2003). *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. RAND Corporation.

[26] List, C. and Pettit, P. (2002). Aggregating sets of judgments: an impossibility result. *Economics and Philosophy*, 18(1), 89–110.

[27] List, C. (2012). The theory of judgment aggregation: an introductory review. *Synthese*, 187(1), 179–207.

[28] Longino, H. E. (2002). *The Fate of Knowledge*. Princeton University Press.

[29] Mastrandrea, M. D., Field, C. B., Stocker, T. F., et al. (2010). *Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties*. Intergovernmental Panel on Climate Change.

[30] McMahan, B., Moore, E., Ramage, D., Hampson, S., and Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of Machine Learning Research*, 54, 1273–1282.

[31] Morris, P. A. (1974). Decision analysis expert use. *Management Science*, 20(9), 1233–1241.

[32] Moss, R. H. and Schneider, S. H. (2000). Uncertainties in the IPCC TAR: recommendations to lead authors for more consistent assessment and reporting. In *Guidance Papers on the Cross Cutting Issues of the Third Assessment Report*, 33–51. World Meteorological Organization.

[33] Nadasdy, P. (1999). The politics of TEK: power and the integration of knowledge. *Arctic Anthropology*, 36(1–2), 1–18.

[34] Nii, H. P. (1986). Blackboard systems, part one: the blackboard model of problem solving and the evolution of blackboard architectures. *AI Magazine*, 7(2), 38–53.

[35] Prakken, H. (2010). An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2), 93–124.

[36] Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press.

[37] Stone, M. (1961). The opinion pool. *Annals of Mathematical Statistics*, 32(4), 1339–1342.

[38] Tebaldi, C. and Knutti, R. (2007). The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions of the Royal Society A*, 365, 2053–2075.

[39] van der Sluijs, J. P. (2005). Uncertainty as a monster in the science-policy interface: four coping strategies. *Water Science and Technology*, 52(6), 87–92.

[40] Walker, W. E., Haasnoot, M., and Kwakkel, J. H. (2013). Adapt or perish: a review of planning approaches for adaptation under deep uncertainty. *Sustainability*, 5(3), 955–979.

[41] Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.

[42] Wooldridge, M. (2009). *An Introduction to MultiAgent Systems*, 2nd edn. Wiley.

[43] Wynne, B. (1992). Uncertainty and environmental learning: reconceiving science and policy in the preventive paradigm. *Global Environmental Change*, 2(2), 111–127.

Draft