ESSAI-2024
Self-Governing Multi-Agent Systems
L3/10: Distributive Justice

Jeremy Pitt and **Asimina Mertzani**
Department of Electrical and Electronic Engineering
Imperial College London

Aims and Objectives

- Aims
  - To understand and analyse *n*-agent social interaction through the lens of computational justice
- Objectives
  - Design and implement algorithms for determining some aspect of 'correctness' in the outcomes of deliberation and decision-making SGMAS

Example

- The *n*-person flat-share biscuit-distribution stand-off
  - You and your flatmates have successfully overcome the kitchen-cleaing stand-off and everyone helped to clean the kitchen
  - One person cleaned the cooker, one person cleaned the fridge, one person did the washing up, one person mopped the floor, ...
  - You all got your utility for for time-loss, convenience and disease-avoidance
  - You decide to celebrate with a biscuit
  - But you find there is only one chocolate Hob-Nob left
- So: who gets the biscuit?

Institutions for Open System

- Use some mutually agreed, conventional **rules/procedures** to ensure that
  - Collective goals are achieved (sustainability)
  - Individual goals are considered as well (satisficing)
  - Balance between all these goals is mutually agreeable
- Is it **fair**?
- Address fairness question through **Distributive Justice**

Distributive Justice: what is it?

- It is concerned with **fairly** allocating goods (also benefits, duties, burdens) to a set of actors in the society
- Aristotle's principle: *"Equals should be treated equally, and unequals unequally, in proportion to the relevant similarities and differences"*
- Three main families of distributive justice theories:
  - *Equality and need*
  - *Utilitarianism and welfare economics*
  - *Equity and desert*

Distributive Justice: what is it?

- Equality and need
  - Concern for the welfare of *those least advantaged* in the society
  - *Need principle*: equal satisfaction of basic needs
  - Some theories: Egalitarianism, Rawl's theory, Marxism
- Utilitarianism and welfare economics
  - Maximising the *global surplus* (outcome, utility, satisfaction)
  - Does not deal with individual outcomes, but in the *aggregation* of these
  - Theories: utilitarianism, Pareto principles, envy-freeness
- Equity and desert
  - *Dependence* of allocations on the actions of each individual
  - *Equity principle*: an individual should receive an allocation that is proportional to her contributions (either positive or negative) to the society
  - Theories: equity, desert and Nozick's theory

Fairness Criteria

- What **fairness criteria** to use to distribute the resources?
  - *Egalitarian*: maximise satisfaction of most disadvantaged agent
  - *Envy-free*: no agent prefers the allocation of any other agent
  - *Proportional*: all agents receive the same share
  - *Equitable*: each agent derives the same utility
  - What about 'Ostrom' principles?
    - Congruence with 'the environment'
    - Those affected by the policy should participate in the selection
    - Especially when g puts others "quite into the shade in point of nastiness"

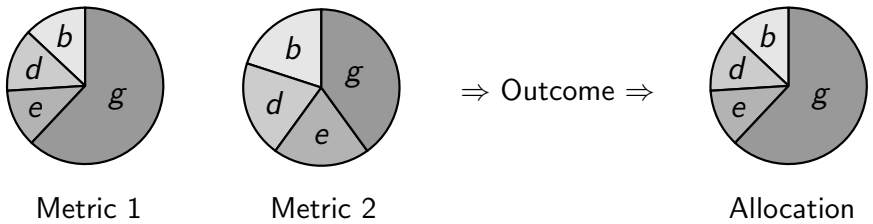

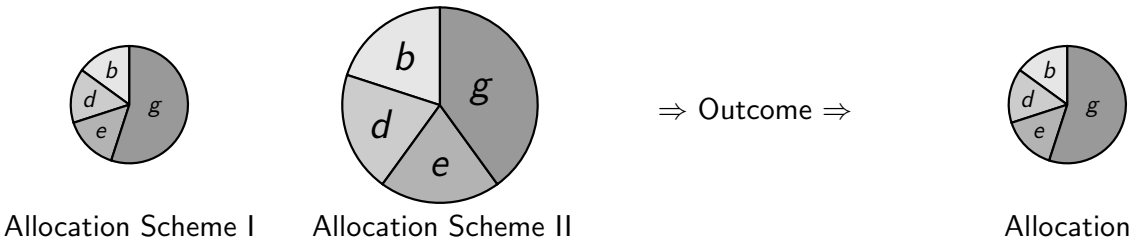

Limitations

- Limitations of existing fairness criteria:
  - Many not appropriate under an economy of scarcity
  - Focus on a single aspect (monistic)
  - Often disregard temporal aspects (e.g. repeated allocations)
- Alternative:
  - Use **multiple** criteria simultaneously (pluralistic)
  - Rescher's theory of Distributive Justice

Rescher's Theory of Distributive Justice

- Rescher proposes: throughout history, seven general principles of distributive justice — to treat people according to...
  - ... as equals
  - ... needs
  - ... actual productive contribution
  - ... efforts and sacrifices
  - ... a valuation of their socially-useful services
  - ... supply and demand
  - ... ability, merit or achievements
- These Rescher called **canons of distributive justice**
- Each canon, in isolation, is inadequate to achieve fairness
- Instead, in context
  - Identify which canons are relevant — these Rescher called **legitimate claims**
  - In case of plurality, decide how to combine them
  - In case of conflict, decide how to reconcile them

Legitimate Claims as Voting Functions

- Each canon C_i treated as a Borda voter on agents
  - It ranks agents according to some features (e.g. needs, contribution...)
  - It assigns a score to each agent, $B_i(a)$
- To combine claims, a weight w_i is attached to each canon
- Final Borda score of agent a is:

$$B(a) = \sum_{i=1}^{n} w_i \cdot B_i(a)$$

- Use final Borda ranking as a queue to allocate resources
- Allocate agents' full requests until no more resources available

Legitimate Claims in action

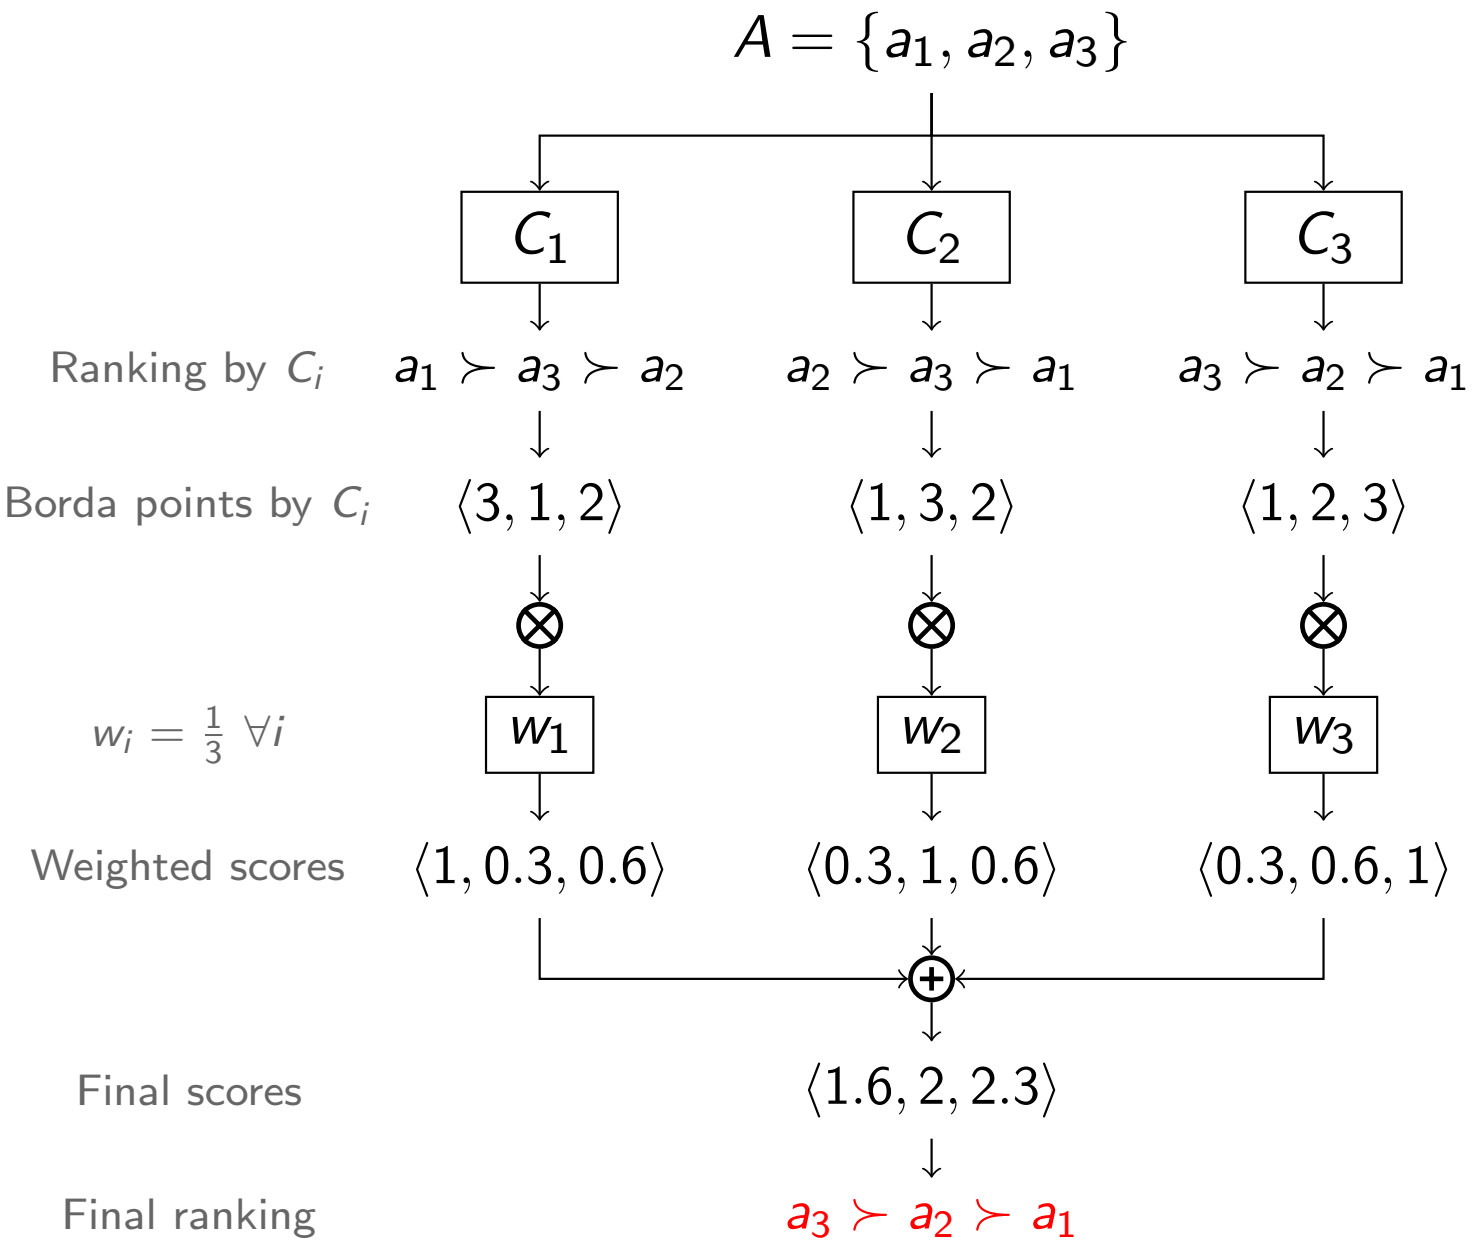

Recall: LPG'

- Agents join a community C to share access to a CPR
  - Defined an institution (set of conventional rules) to regulate how to make provision to/appropriation from the resource
  - Now define operational choice rules to determine 'priority' order for access to resource
- Focus on observable actions

$d_i(t)$	Demand of ...	...agent i at time t
$p_i(t)$	Provision of ...	
$r_i(t)$	Allocation to ...	
$r'_i(t)$	Appropriation of ...	
$role\_of(i,t)$	Role of ...	
$\mathbf{T}_{\{i \in C\}}$	Rounds agent i present in community C	

Legitimate claims in LPG'

Equals	Average allocation ↑	$\frac{\sum_{t=0}^{T} r_i(t)}{T}$
	Allocation frequency ↑	$\frac{\sum_{t=0}^{T} (r_i(t)>0)}{T}$
	'Satisfaction' ↑	$\sigma_{i,C}(t+1) = \ldots$
Needs	Average demands ↑	$\frac{\sum_{t=0}^{T} d_i(t)}{T}$
Contribution	Average provision ↓	$\frac{\sum_{t=0}^{T} p_i(t)}{T}$
Effort	Number of rounds present ↓	$\lvert \mathbf{T}_{\{i \in C\}} \rvert$
Social utility	Time as *head* ↓	$\lvert \{t \lvert role\_of(i,t) = head\} \rvert$
Supply & demand	Compliance ↓	$\lvert \{t \lvert r'_i(t) = r_i(t)\} \rvert$
Ability, merits...	n/a	

Self-Determining the Weights (1)

- Instead of fixing the weights of each canon, allow the agents to modify them
- At the end of each round
  - Agents vote for the canons in order of preference (according to rank given by each canon) using a modified Borda count
    - Allows for some candidates to have the same number of points (equal preference)
  - Borda score computed for each canon
  - Canons with better than average Borda score have weight increased, otherwise decreased

- This reflects Ostrom's Principle 3: *"those affected by the operational-choice rules participate in the selection and modification of those rules"*

Self-Determining the Weights (2)

Determining the canons' weights

	Points given by			Ranking	Points given to		
	C_1	C_2	C_3		C_1	C_2	C_3
a_1	3	1	1	$\langle C_1, C_2 \sim C_3 \rangle$	3	1.5	1.5
a_2	1	3	2	$\langle C_2, C_3, C_1 \rangle$	1	3	2
a_3	2	2	3	$\langle C_3, C_1 \sim C_2 \rangle$	1.5	1.5	3
					5.5	6	6.5

$$\text{Average Borda score} = 6 \Longrightarrow \begin{cases} w_1 & \downarrow \\ w_2 & = \\ w_3 & \uparrow \end{cases}$$

Some Experiments

- Experimental platform using PreSage-2 [Macbeth et al, 2012]
- Independent Variables: agent population
  - Number of agents, % of non-compliant agents
  - Propensity to cheat on provision and/or appropriation
  - Initial satisfaction, dissatisfaction threshold
  - Coefficients a, b, c (utility), α, β (satisfaction), γ (autonomic mechanism)
- Dependent variables
  - Utility of the compliant/non-compliant agents
  - Endurance of compliant/non-compliant agents
  - Fairness measure: Gini inequality index
- Questions
  - Single cluster: effect of plurality, and self-organisation
  - Single cluster: Comparison with alternative allocation policy
  - Multiple cluster: effect of allocation method on cluster preference

Some Results

- Compare self-organising legitimate claims, fixed weights, random and ration allocation methods
- Self-organising legitimate claims...
  - ...was the only method producing endurance of the system and benefiting compliant agents
  - ...was the fairest* method (wrt to ration and fixed LC)
  - ...was preferred by the compliant agents
  - ....leads to a very fair overall allocation by doing a series of rather unfair allocations

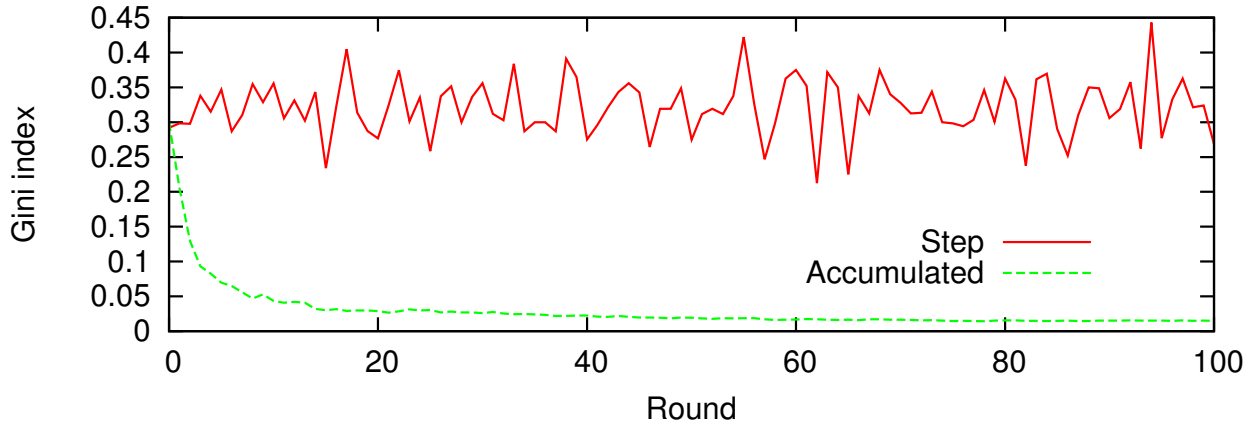

*Using Gini inequality index over accumulated allocations to measure fairness

Key Features of Open Systems

- **Self-determination**
  - Rules for resource allocation and how to choose them determined by the entities themselves
- **Expectation of error**
  - Rehaviour contrary to specification should be expected (be it by accident, necessity or malice)
- **Enforcement**
  - Sanctions for non-compliance should be implemented
- **Economy of scarcity**
  - Eufficent resources to keep appropriators satisfied at the long-term, but insufficient to meet all demands at a particular time-point
- **Endogeneous resources**
  - Computing the allocation must be 'paid for' from the same resources being allocated
- **No full disclosure**
  - Appropriators are autonomous and their internal states cannot be checked

Computational Justice

- It is sustainable. It is fair.
- But what about
  - Is the allocation method **effective**? Is it **efficient**? Is it **fit-for-purpose**?
  - Are decision makers **accountable**?
  - Do those affected by the rules **participate** in their selection?
  - Are **punishments** for non-compliance proportional to the severity of the offence?
- Answering all these questions requires a framework for **computational justice**

Reminder: Key Features of SGMAS

- **Self-determination**
  - Rules for resource allocation and how to choose them determined by the entities themselves
- **Expectation of error**
  - Rehaviour contrary to specification should be expected (be it by accident, necessity or malice)
- **Enforcement**
  - Sanctions for non-compliance should be implemented
- **Economy of scarcity**
  - Satisfaction vs. Satisficing
- **Endogeneous resources**
  - Computing the allocation must be 'paid for' from the same resources being allocated
- **No full disclosure**
  - Appropriators are autonomous and their internal states cannot be checked

Computational Justice and SGMAS

Key features	*Justice*
Self-determination	**Natural**
Expectation of error	**Retributive**
Enforcement	
Economy of scarcity	**Distributive**
Endogeneous resources	**Procedural**
No full disclosure	**Interactional**

Computational Justice and SGMAS

Key features		*Justice*
Self-determination	participation, inclusion, voting (1)	**Natural**
Expectation of error		**Retributive**
Enforcement		
Economy of scarcity		**Distributive**
Endogeneous resources		**Procedural**
No full disclosure		**Interactional**

(1) Pitt et al, *Coordination, conventions and the self-organisation of sustainable institutions*. PRIMA 2011

Computational Justice and SGMAS

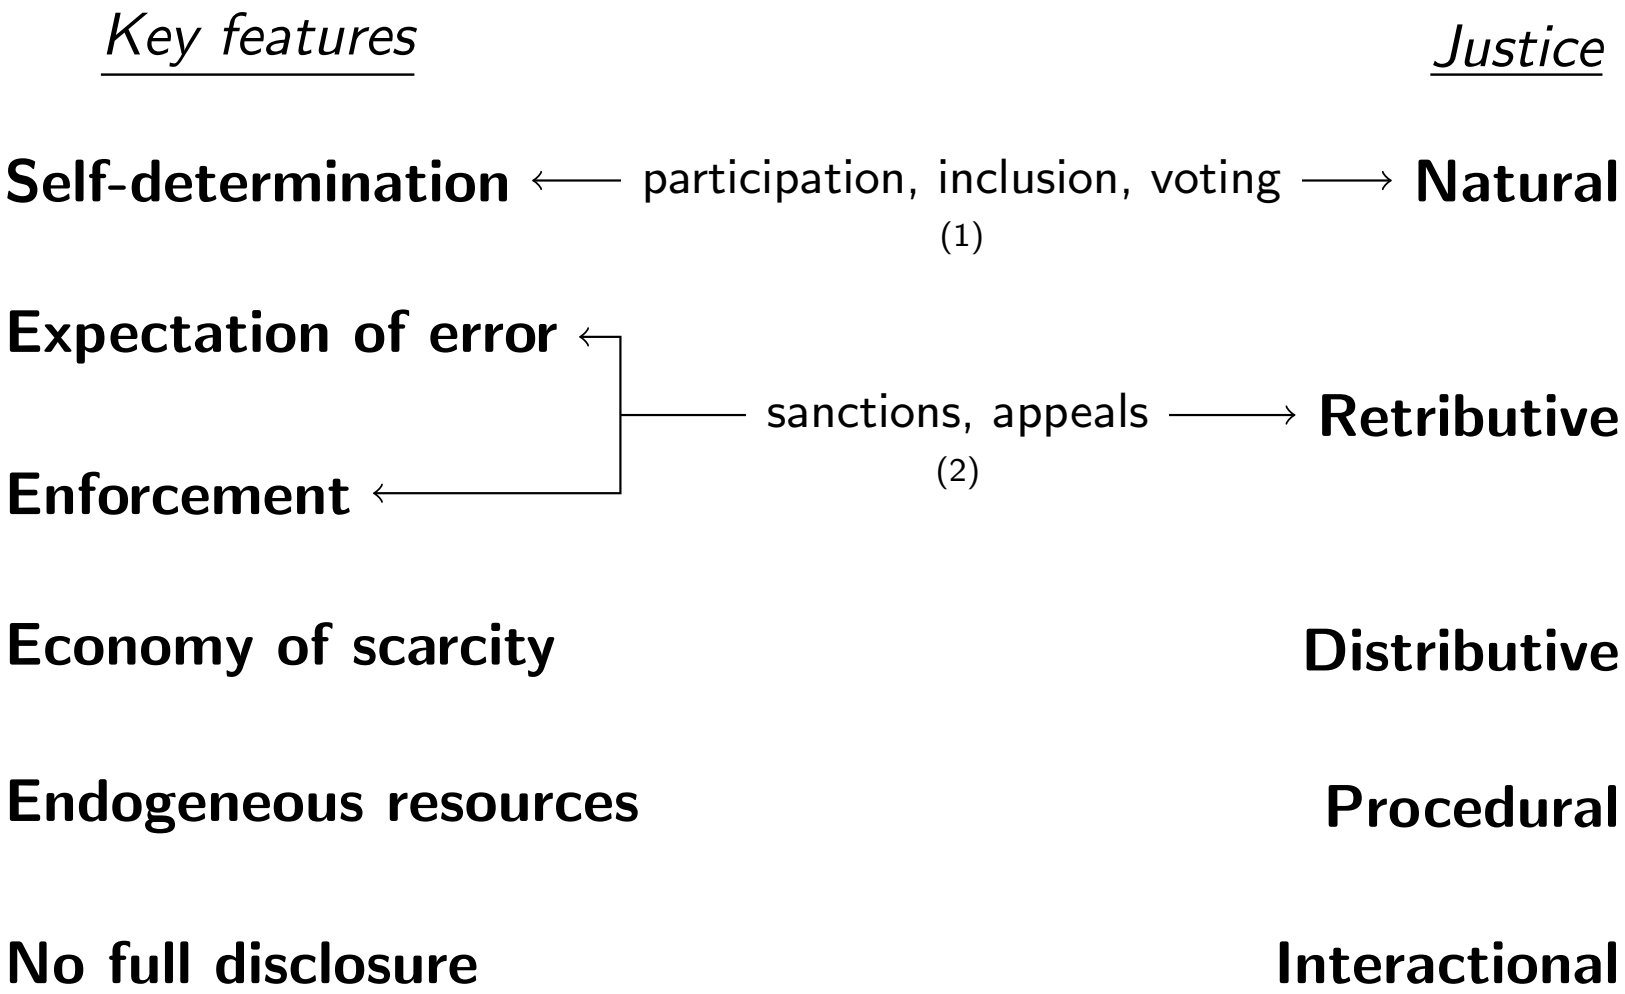

(1) Pitt et al, *Coordination, conventions and the self-organisation of sustainable institutions*. PRIMA 2011
(2) ____, *Provision and appropriation of common-pool resources without full disclosure*, PRIMA 2012

Computational Justice and SGMAS

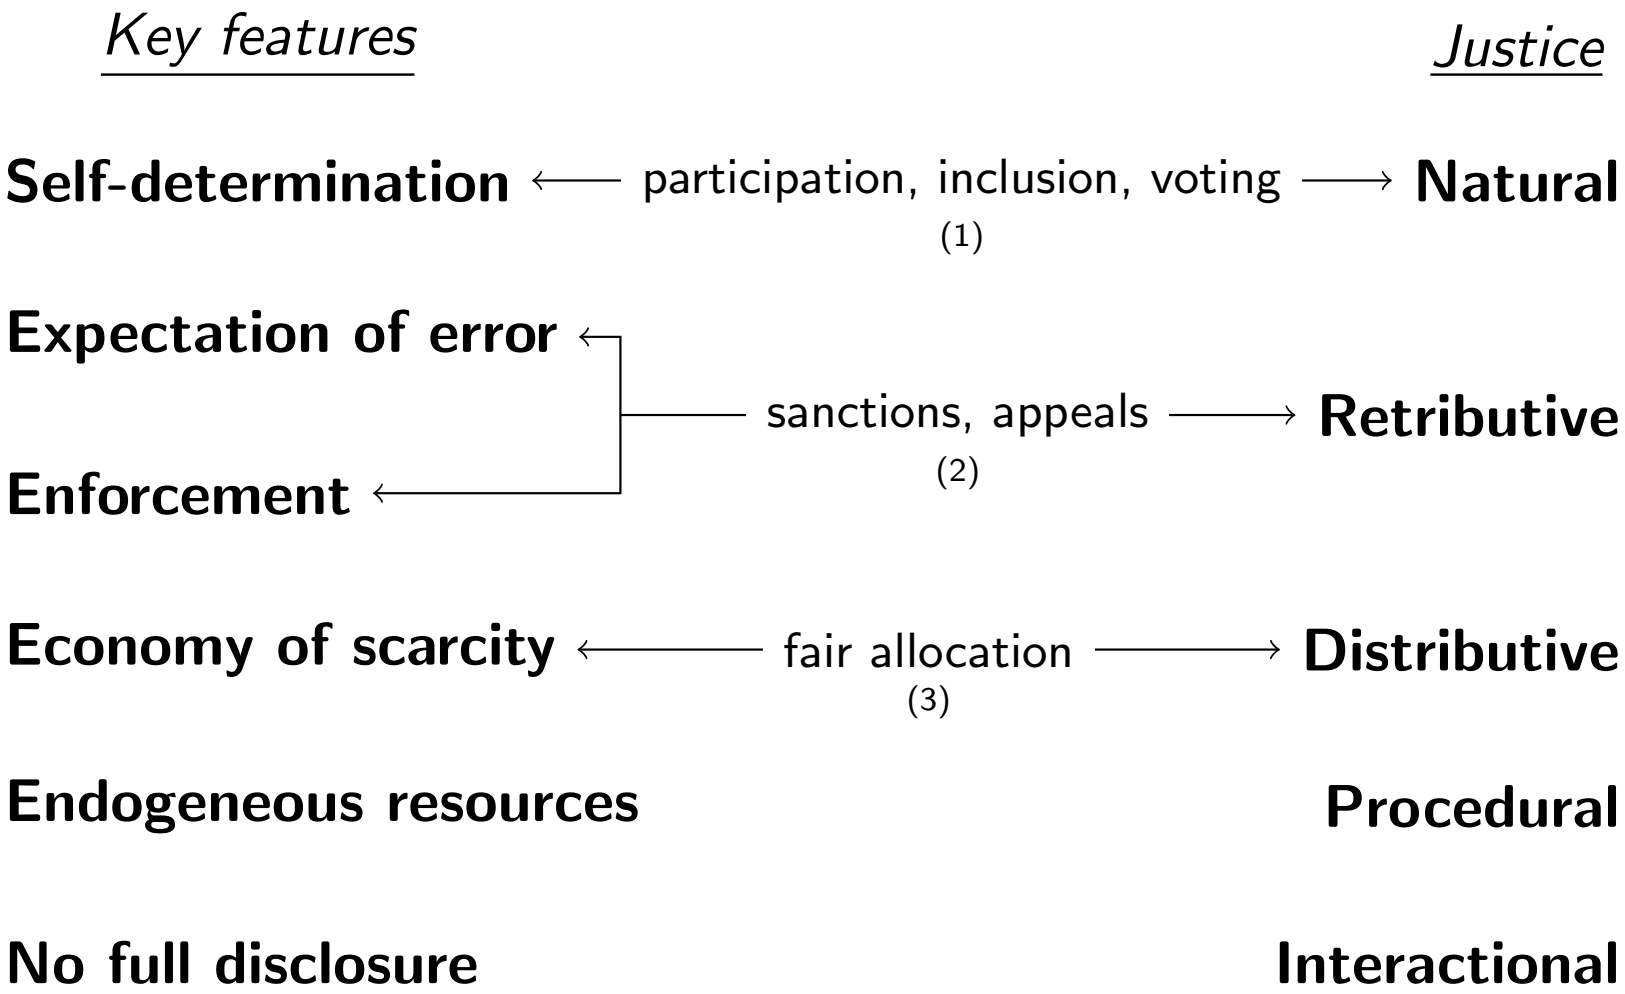

(1) Pitt et al, *Coordination, conventions and the self-organisation of sustainable institutions*. PRIMA 2011
(2) _____, *Provision and appropriation of common-pool resources without full disclosure*, PRIMA 2012
(3) _____, *Self-organising common-pool resource allocation and canons of distributive justice*, SASO 2012

Computational Justice and SGMAS

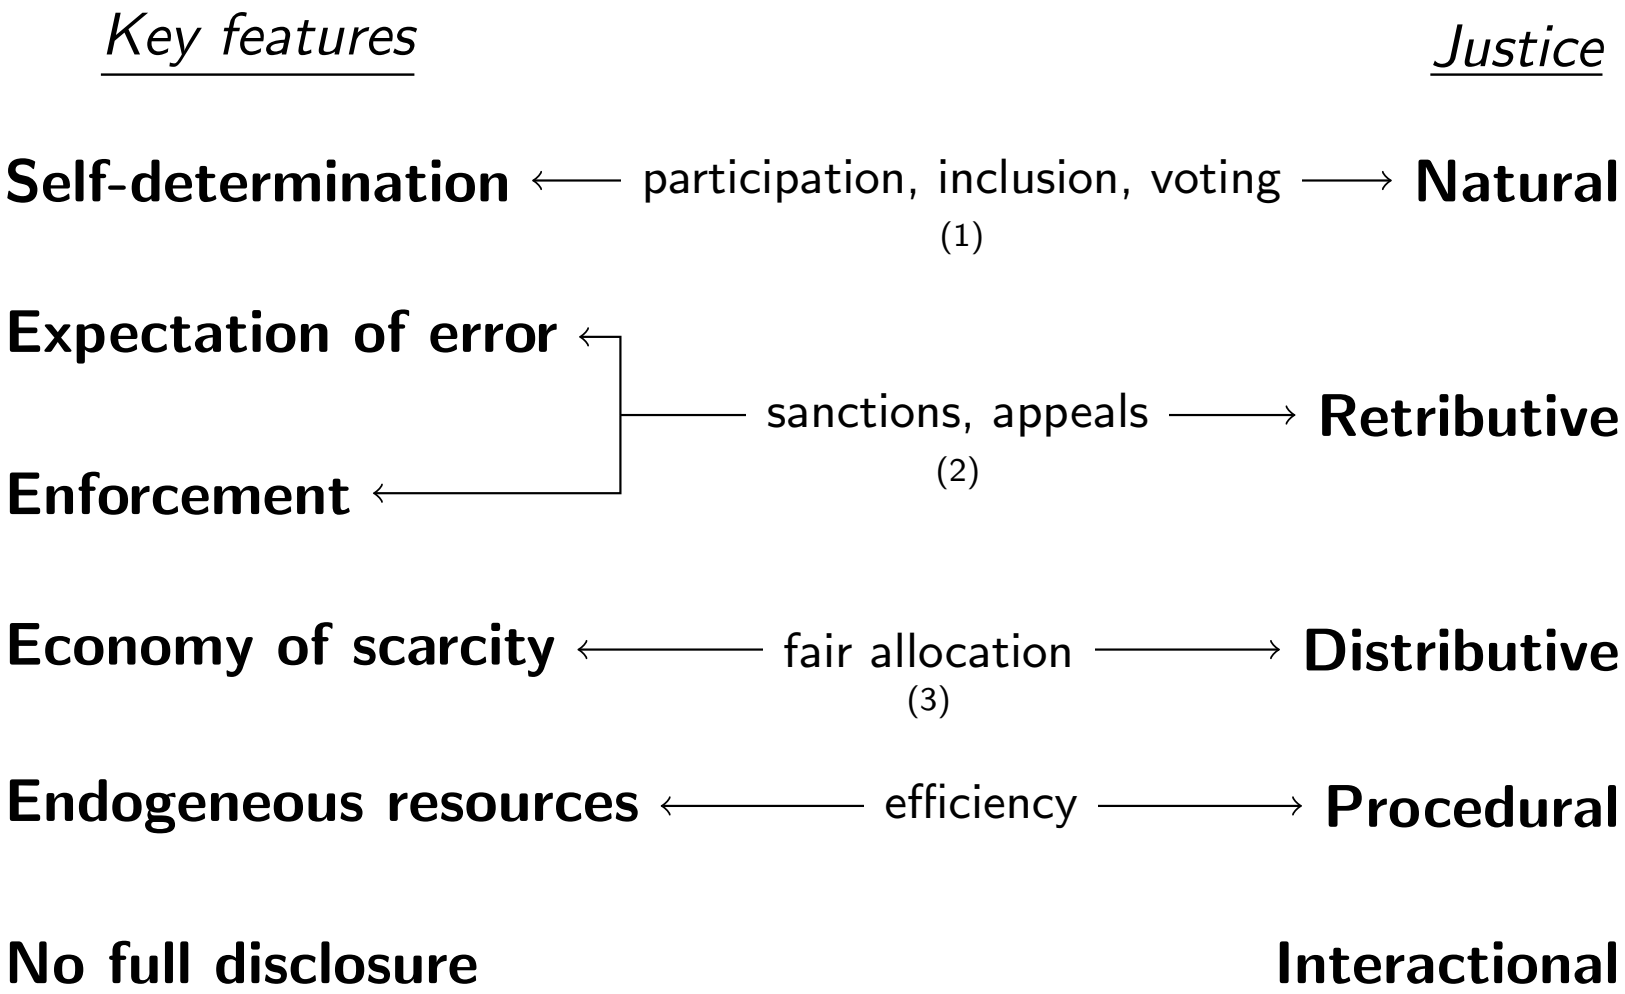

(1) Pitt et al, *Coordination, conventions and the self-organisation of sustainable institutions*. PRIMA 2011
(2) _____, *Provision and appropriation of common-pool resources without full disclosure*, PRIMA 2012
(3) _____, *Self-organising common-pool resource allocation and canons of distributive justice*, SASO 2012

Computational Justice and SGMAS

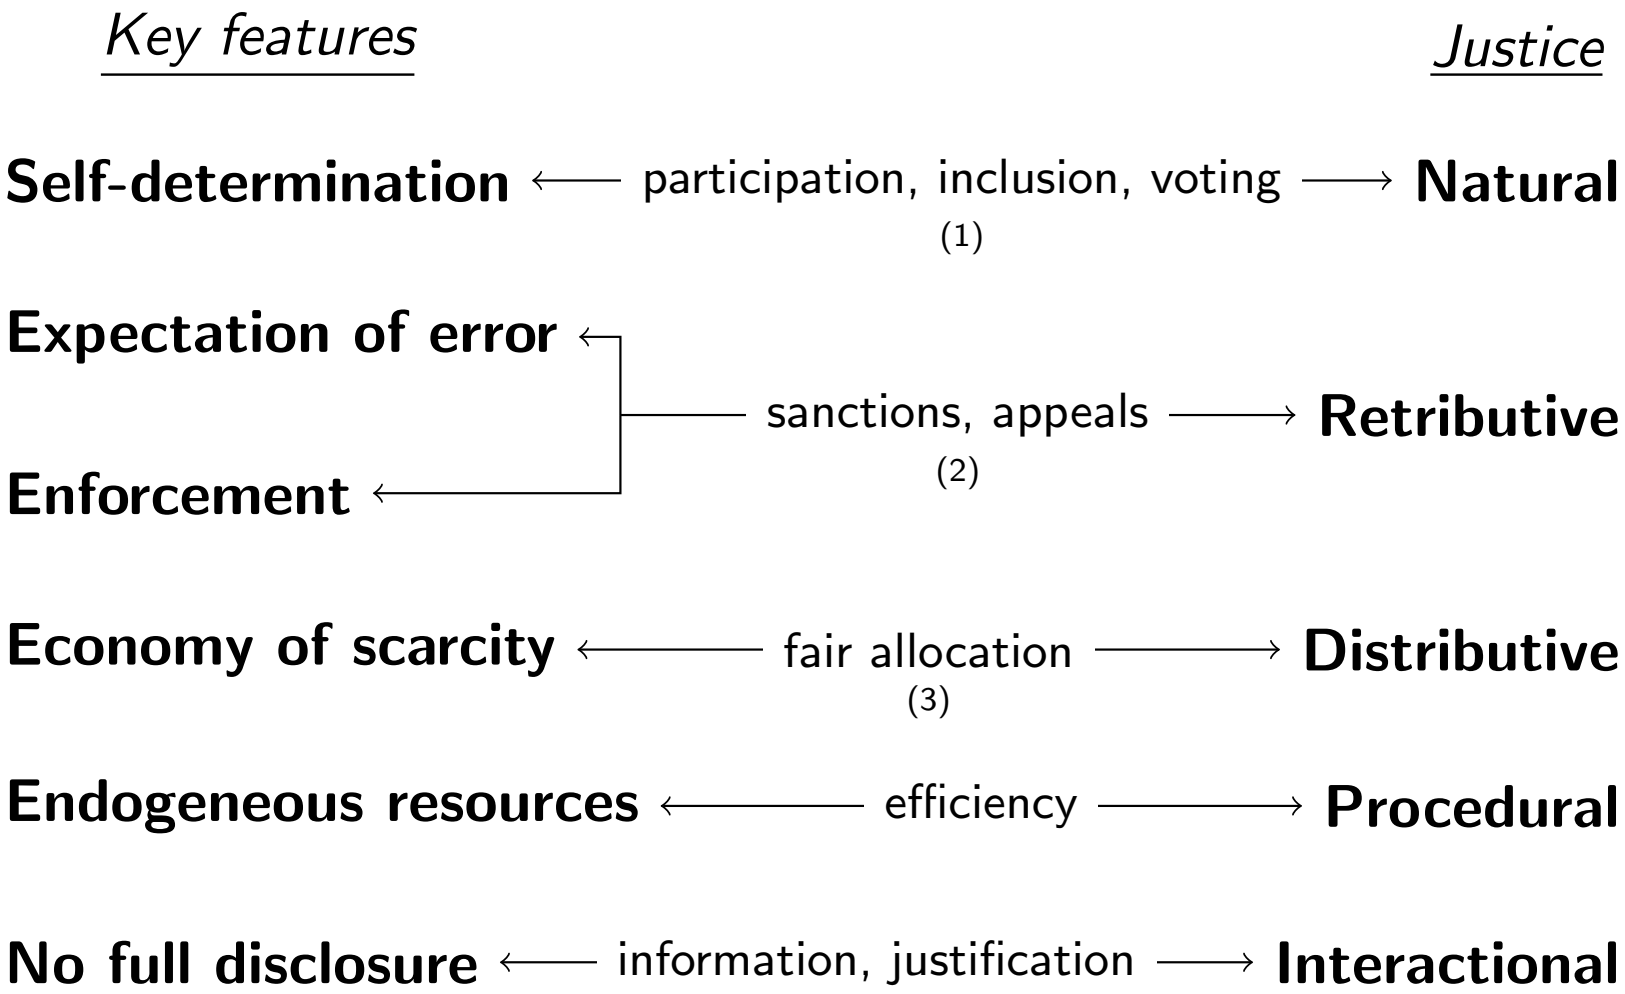

(1) Pitt et al, *Coordination, conventions and the self-organisation of sustainable institutions*. PRIMA 2011
(2) _____, *Provision and appropriation of common-pool resources without full disclosure*, PRIMA 2012
(3) _____, *Self-organising common-pool resource allocation and canons of distributive justice*, SASO 2012

Summary and Conclusions

- Should you ever hear: "there is no fair way to do it"
- You can say: "yes there is"
- Presented an algorithm for distributive justice
  - Objectively, if everyone agrees on the legitimate claims and their adaptation
- But there is still more to do
- Fairness is also a subjective experience...